



6

Elements of Fundamental Reform

In this chapter, we examine the common elements of many proposals for fundamental tax reform. We begin this exercise by examining the appeal of so-called *flat* taxes.

Many reform proposals that differ tremendously from each other bill themselves as flat taxes. That flatness has been elevated to the highest compliment to be paid to a tax structure is a bit surprising, if only because in other contexts “flat” is not always a good thing. Just think of beer, musical notes, or tires. *Webster’s Unabridged Dictionary* lists sixteen major definitions for *flat* as an adjective, ranging from praiseworthy ones such as “level” and “exact” to much less attractive ones, including “shallow,” “dull and stupid,” “commercially inactive,” and even “having no money.” In the context of taxation, the word *flat* conjures up *Webster’s* definition of “not varying.” But even that is open to multiple interpretations; most of us wouldn’t want a tax in which liability does not vary with circumstances such as income or wealth.

There are three distinct dimensions of flatness in a tax—a single tax rate, a consumption tax base, and a clean tax base. The single rate is what most people pick up on, but the other two dimensions would represent even more fundamental changes in the way we tax ourselves. Some proposals seek flattening in all three dimensions; other plans focus on only one or two aspects of flatness. This chapter explores each of these aspects of flatness.

A Single Rate

The most eye-catching feature of flat taxes is the flat *rate*. In place of our current system of graduated tax rates that increase with higher incomes, all or most taxpayers under a flat tax system would be subject to a single rate of tax. Although a truly flat-rate tax would apply the



single rate to the entire tax base, from the first dollar to the last, in most flat tax proposals the single rate applies only to the base in excess of some exemption level. For that reason, they really are a form of graduated tax, with an initial bracket to which a zero tax rate applies, plus an open-ended bracket subject to a single tax rate. Under such a system, the ratio of tax liability to the tax base (i.e., the average tax rate) is zero until the exemption level is reached and then increases gradually with the tax base until, for very high incomes, the average rate is nearly the single rate. Moreover, the degree of progressivity can be varied by adjusting the level of tax-exempt income and the tax rate. A flat-rate tax that exempts all income below \$100,000 and levies a 50 percent tax on income above \$100,000 is considerably more progressive than a flat tax rate that exempts only \$10,000 of income and levies a 20 percent tax above that.

To distinguish this aspect of flatness from the others, we refer to a tax system with one tax rate as a *single-rate tax*. Replacing the graduated tax rates with a single rate can be accomplished independently of any and all of the changes in the tax base, to be discussed below, that are usually associated with flat taxes. A single rate can be applied to a narrow, preference-ridden base or to a broad, clean base, and it can be applied to all of income or just to the portion that is consumed. Similarly, we can certainly clean up the tax base while maintaining graduated rates and implement a consumption tax in a way that allows us to preserve graduated rates.

If nothing else changed, replacing the graduated tax rate structure with a single tax rate would make the distribution of tax liabilities dramatically less progressive. Figure 6.1 illustrates this point. The thick black line depicts how average income tax rates vary with adjusted gross income in 2005 for typical married couples with two children who have average deductions for their income levels. One thing to note about current law is that income tax rates are currently *negative* for low-income families, due mainly to the earned income tax credit. A second is that the tax-exempt threshold is already quite large: in 2005, a couple with two children did not face any positive income tax liability until its income reached just above \$41,000. Third, above that tax-exempt threshold, income tax burdens were still distributed very progressively, as *marginal* rates in 2005 gradually rose from 10 percent at the lowest levels of taxable income up to 35 percent above \$326,450. As a result, *average* tax rates continue to rise well into the upper end of the income distribution.¹

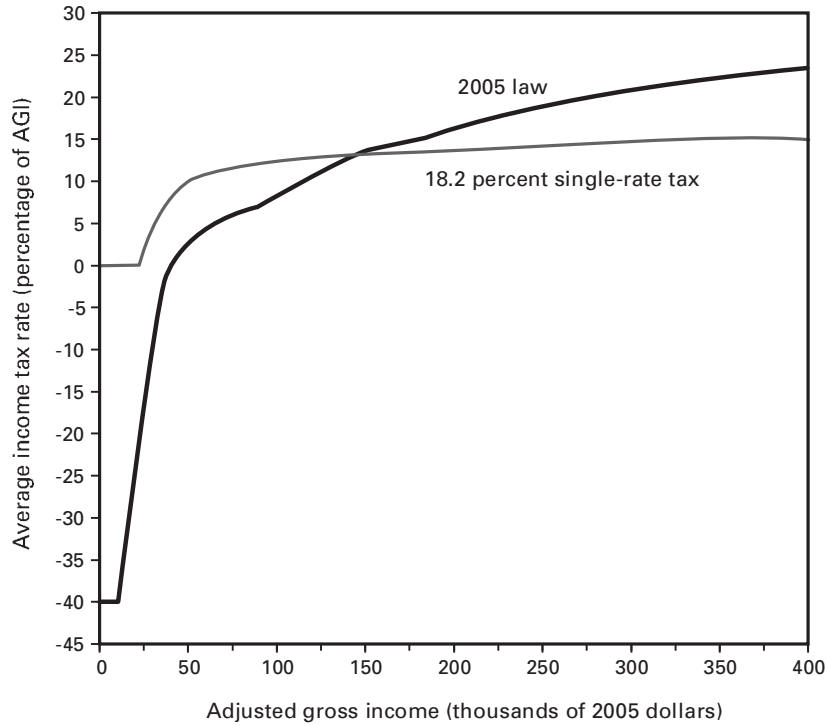


Figure 6.1
Effect of a revenue-neutral elimination of graduated tax rates, credits, and the AMT on average personal income tax rates for typical married couples with two children in 2005.
Source: Authors' calculations based on data from Balkovic (2007) and IRS, *Individual Income Tax Returns 2004* (2006a).

The thinner line in figure 6.1 illustrates the effects of switching to a single rate with a large tax-exempt threshold, in the spirit of some flat tax proposals floated in recent years. To isolate the effects of switching to a single rate from the effects of other components of tax reform, we remove the main features (other than personal exemptions and the standard deduction) that contribute to the current system's progressivity but otherwise leave the tax base the same. Thus, the graduated rate structure, tax credits, and the alternative minimum tax are removed, but all the existing deductions, exemptions, and exclusions are retained. We calculate that to raise the same revenue as the income tax did in 2005, such a tax would require a 18.2 percent tax rate.² The standard deduction and personal exemption would continue to provide a large



tax-exempt threshold (\$22,800) for these families. As advocates of flat taxes emphasize, as long as there is an exempt level of income, such a flat-rate tax is indeed still somewhat progressive: average tax rates do rise with income. However, such a tax is significantly less progressive than the current system. Abandoning the EITC greatly increases average tax rates for the lowest-income families, while raising the 10 percent and 15 percent rates to 18.2 percent and eliminating child credits increases tax burdens for most of the middle class. At upper income levels, average tax rates start to flatten out much earlier and settle at a level lower than under current law. Families with AGI below \$144,000 face tax increases, and those with AGI above that level receive tax cuts. A family with \$350,000 of adjusted gross income (AGI) would get a tax cut of \$25,484, or about 7.3 percent of AGI. A family with \$75,000 of AGI, by contrast, faces a tax increase of \$4,268, or 5.7 percent of AGI.

This exercise points out starkly how much the appeal of a single-rate tax depends on exactly what rate is enacted and to what extent the rate change is accompanied by measures that broaden the tax base. In practice, most advocates of a flat tax are also advocates of a flat, *low* tax rate, usually in the high teens or low twenties. As shown above, switching to such a rate by itself would radically shift tax burdens away from upper-income families and toward low- and middle-income families. Therefore, cutting back on exclusions and itemized deductions that disproportionately benefit upper-income families would be absolutely essential if we want to raise the same revenue as the current tax system with a single low rate but at the same time mitigate the negative impact on low- and moderate-income families. Even with substantial base broadening, however, it would be practically impossible to design a tax system with a single rate that is significantly below the current top rate of 35 percent that replicates the current tax liabilities on people with extremely high incomes. As a result, somebody else will have to pay more in taxes.

As we noted in chapter 3, survey evidence suggests that the flat tax apparently appeals to some people because they think it will make the rich pay *more* in taxes. This is almost certainly not true. Abandoning graduated rates for a single rate will lower the tax burdens on very high-income people, even if we get rid of every loophole and every bit of evasion there is.

So why go to a single rate? Perhaps the most commonly cited reason is to improve economic incentives. Raising the same amount of revenue



with a graduated rate structure requires that, on average, people face higher *marginal* tax rates, which, as discussed in chapter 4, can discourage work and saving and cause a whole host of other economic distortions and inefficiencies. Thus, tax policy design must confront a trade-off between progressivity and economic prosperity. In principle, improved economic growth could eventually compensate some of the people who initially face higher tax burdens as a result of switching to a single rate. As shown in chapter 4, however, the economic costs of progressivity are uncertain and are almost certainly not as high as they are often made out to be by many political advocates of a flat rate. Other things being equal, lower marginal rates are better for the economy, but economics reveals nothing magical about a single rate. For example, the economic cost of having two low marginal rates—say, 15 percent and 25 percent—is not likely to be significantly higher than the cost of a single 20 percent rate.

Some argue that a single-rate tax structure is “neutral” toward distributional issues and thereby avoids “class warfare.” By advocating a single rate, they suggest they are transcending the controversy over fairness and are promoting the system that maximizes economic performance. Neither of these arguments is correct. A flat-rate tax with a low rate generates a distribution of tax burden that is less progressive than the current graduated tax rate structure, but it certainly affects the distribution of income and is not *neutral* in any meaningful sense. The only tax system that would truly eliminate all the economic costs of taxation is a lump-sum tax (that is, a poll tax), under which tax *liability*, not the tax *rate*, is the same for everyone, rich or poor. In place of income taxes, we could have a fixed annual charge of \$5,600 per adult, whether that adult is Bill Gates, a homeless person, or yourself.³ This eliminates any tax penalty tied to work, saving, or investing. Presumably, the reason that the proponents of a single-rate tax prefer it to a lump-sum tax is that they find the latter abjectly unfair (or else they are too timid to admit the opposite). Thus, the single-rate tax structure already reflects a balancing between equity and efficiency. A two-rate or multiple-rate structure reflects another balancing, as does a lump-sum tax. We can argue about how best to make this balancing, but no tax system avoids this trade-off, and no tax system is distributionally neutral.

When it comes to simplicity and enforceability, however, there is indeed something special about a single rate, although it’s not what most people think. A single rate doesn’t really simplify your tax return;



in fact, applying the tax rate schedule to your taxable income to calculate how much you owe is the least complicated part of the whole taxpaying process. Most people can just look it up in a table, in which case it hardly matters at all whether there are one, two, or twenty rates. For those (individuals or tax professionals) who use software, the number of different rates certainly makes no difference in the difficulty in tax filing.

The real advantage of a single-rate tax structure is that it potentially facilitates much simpler, business-based, systems of collecting taxes, such as a value-added tax, under which no individuals need file tax returns or remit taxes to the tax authority. Even under an income tax with individual filing, a single rate makes it easier to rely on employers withholding and remitting taxes on behalf of employees, banks withholding and remitting on behalf of those who receive interest, and so on, a major simplification and enforcement advantage.

Reducing the disparity of tax rates also reduces the incentives for individuals to shift taxable income from high-rate to low-rate taxable entities and from high-tax to low-tax periods. Moreover, it minimizes the horizontally inequitable treatment accorded to those with fluctuating incomes. Under a graduated income tax, comparing two families with the same lifetime income, the one with fluctuating income will pay more total tax than a family with a more stable income flow; a flat-rate system eliminates this disparity.

Although these are important advantages, it's debatable whether they offset the distributional consequences of a low flat rate. Moreover, adding another rate on top of the main flat rate doesn't entirely destroy the simplicity advantages. The United Kingdom greatly simplified its income tax system by moving to a single rate for most people and improving its withholding system but still retained a higher marginal rate for a small percentage of high-income individuals.⁴ And having the same *top* rate across individuals, businesses, and different types of income eliminates many of the opportunities for complex income-shifting and tax avoidance schemes, even if there are graduated rates below that top rate.

A Consumption Base

The second element of many fundamental tax reform plans is to change the system so that tax liability is triggered by consumption rather than by income. Admittedly, noneconomists do not associate this aspect of



a tax system with flatness. But to economists, a consumption tax imposes a uniform (call it flat, if you like) tax on current consumption and future consumption, in the sense that it does not alter the reward for postponing consumption (that is, saving). In contrast, an income tax, because it taxes the return to saving, makes consumption in the future more expensive than consumption now. We explain further below.

What Is a Consumption Tax?

For some readers, the word *consumption* may conjure up memories of their introductory college economics class, where they may have first encountered consumption as being something other than an old-fashioned word for tuberculosis. Consumption is just economists' language for what people do when they use up (i.e., consume) goods and services. A consumption tax simply means that the "tax base" (what triggers tax liability) is consumption, as opposed to income, wealth, or some other concept. There are multiple ways to measure consumption. Because annual income—for a family or for the country as a whole—is equal to annual consumption plus saving, consumption can be either measured directly or, alternatively, measured by subtracting saving from income. It is also true by definition that the amount of saving done by a country equals the amount of investment on which its residents have a claim to the returns, and the total income of a country equals the total value of output of factors of production (labor and capital) owned by residents of the country. So another way of measuring consumption is output minus investment. Note that the term *investment* here refers to, as it generally does in economics, a business purchase of a capital good such as a plant or equipment and not to an individual purchasing a share of General Electric stock.

Varieties of Consumption Tax

To make our discussion more concrete, let's consider some examples of consumption taxes. The *retail sales tax* is probably the one most familiar to Americans, as almost all states already operate one. Under this approach, tax liability is triggered by the final sales of goods and services at the retail level, and the tax is collected entirely from retailers (that is, businesses that sell directly to consumers).

Another kind of consumption tax that is much more commonly used around the world, partly because it is easier to administer and enforce, is the *value-added tax* (VAT). Instead of collecting tax only from retailers,



under a VAT businesses at each stage of the production and distribution process must remit tax. Each firm pays tax on its “value added,” which is simply its total sales minus the cost of the inputs it purchases from other businesses. For example, a furniture-building company would remit tax equal to the VAT rate multiplied by its sales of furniture net of the costs of wood, glue, work benches, electricity, and other inputs that it bought for the business. The furniture *retailer* remits tax on its sales to consumers but can deduct the value of its purchases from the furniture *maker* and so on throughout the production and distribution chain, so that all business-to-business purchases ultimately cancel each other out, leaving final sales to consumers as the ultimate tax base.

A close relative to the value-added tax is what economists refer to as “the” flat tax. Businesses owe tax on sales minus inputs purchased, just like under a VAT, but they can also deduct payments to workers. Critically, workers owe tax on their labor income at the same rate of tax as do businesses. Thus, this third approach to taxing consumption is essentially equivalent to a VAT—the only difference being that labor compensation is deducted from the business tax base and then taxed separately, and at the same rate, at the personal level. This method makes it easier to introduce some progressivity by incorporating a large exemption in the personal tax on labor compensation. The Hall-Rabushka flat tax proposal, most recently championed by 1996 and 2000 Republican presidential candidate Steve Forbes and former House majority leader Dick Armey, is the most famous example of this type of tax.

Finally, under a *personal consumption tax*, each year a taxpayer would calculate and report both income and net saving, and the tax liability is based on the difference between the two—that is, income minus saving, which equals consumption. This kind of tax could easily accommodate progressive rates. The major change to the current system that would get us to this tax would be to create unlimited Individual Retirement Accounts, although other changes (such as eliminating interest deductions and accounting for all borrowing and withdrawals in the measurement of net saving) would be necessary as well to avoid massive and destructive loopholes. The most prominent example of a personal consumption tax prototype is the Unlimited Savings Allowance (USA) tax, a plan championed in the mid-1990s by Senator Pete Domenici (R-NM) and former Senator Sam Nunn (D-GA) that never achieved any political traction.



The various methods of taxing consumption differ in how they *appear* to divide the tax burden up among consumers, workers, and businesses because the official tax liability is very different and because different parties write checks to the IRS. As explained in chapter 3, however, except for implementation issues, these differences are irrelevant to who ultimately bears the burden of the tax and to what are its economic effects. Thus, the different approaches have more in common than it might appear on the surface. In the next chapter, we examine in more detail how each of these approaches would work, explore the differences among them, and evaluate them. In this section, we focus on the many elements that are common to all consumption taxes.

Consumption Taxes, Income Taxes, and Incentives to Save and Invest

The key distinction between an income tax and a consumption tax is that income taxes reduce the incentive to save and invest, whereas consumption taxes do not. The easiest way to see this is to consider an example. Suppose you get a \$100 bonus at work and are trying to decide whether to spend it today or to put it in the bank and save it for next year at a 10 percent interest rate. If there were no taxes, you'd have a choice between consuming the \$100 today or saving it so you have \$110 next year, a 10 percent reward for saving.

Now let's see how things change when there's a single-rate 20 percent income tax. If you choose to spend your raise immediately, you'll get to consume \$80—your take-home pay after the income tax. If you decide instead to save it, your \$80 will earn an extra \$8 in interest. An income tax, however, subjects that interest to taxation as well. After taxes, you only get \$6.40 in interest—so your choice is between consuming \$80 today or \$86.40 next year. The income tax not only reduces your take-home pay but also reduces the reward to saving from 10 percent to 8 percent.

Now consider what happens when, instead of an income tax, there is a retail sales tax levied at a rate of 20 percent of the value of purchases (inclusive of the tax).⁵ If you spend your \$100 raise today, you are subject to a 20 percent tax immediately, so you get to consume \$80 after taxes, just like you would under the income tax. What if you prefer to save the bonus so you can buy goods and services next year? Then you put the \$100 bonus in the bank, and because it earns 10 percent interest, you'll have \$110 in your account next year: —note that, unlike under an income tax, no tax is triggered by the interest



earned. When you withdraw your money and spend it, you pay the 20 percent sales tax, leaving you with 80 percent of \$110, or \$88, of after-tax consumption. Under the consumption tax, your choice is thus between consuming \$80 today and consuming \$88 in a year. The reward for saving is still 10 percent, exactly what it would be in the absence of any tax at all.

This is what economists mean when they say that a consumption tax is “neutral” between current and future consumption: the terms of the choice between consumption today and consumption in the future are the same as they would be in the absence of taxes. In our example, you get the 10 percent reward under the consumption tax, just as you would if there were no taxes, while the income tax leaves you with only an 8 percent reward for saving. Our example is admittedly simplified, but complicating it in any number of ways, such as considering multiple time periods or allowing inflation, would not change the basic result. Moreover, the argument applies to any form of consumption tax, including not only retail sales taxes but also value-added taxes, the flat tax, and a personal consumption tax.

Some of these tax systems require businesses to remit taxes, so to show that under a consumption tax there is no reduction in the return to saving *or investing*, we need to investigate further. It turns out that there is a key difference between how business investment is treated under the two systems. An income tax allows deductions for capital goods as they *depreciate* (i.e., lose their value due to wearing out or technological obsolescence), while a consumption tax allows the cost of capital goods to be deducted from the tax base in full immediately, which is known as *expensing*. Modifying our example a bit can show why expensing removes the tax system’s impact on the incentive of business to invest in the same way that a consumption tax has no impact on the return to saving.

Suppose you’re the owner of a small business that has just earned a \$100 profit. You could take the portion that is left over after taxes out of the business today and use it to buy a new DVD player. Alternatively, you could invest the profits in the business by purchasing a new piece of equipment that will enable the business next year to produce goods that are worth 10 percent over and above the cost of the investment. You know that immediately after producing the goods, the equipment will break down and produce nothing more.

Under either an income or a consumption type of business tax, if you pay the profit to yourself today and consume it, it will be subject to



tax. If the tax rate is 20 percent, you have the option of consuming \$80 after tax today. The difference is in what happens if you invest in the equipment. Under the income tax, you would still pay a 20 percent tax on your \$100 profit today, even if you were to invest in the equipment; there is no deduction for the cost of the equipment because it hasn't worn out yet. So you have \$80 left over to invest, which will produce \$88 in goods for you next year. Next year, you will have \$8 in taxable income—\$88 in sales, minus an \$80 deduction for depreciation. You would pay 20 percent of taxable income, or \$1.60, in tax next year, leaving you with \$86.40 to consume. So the business income tax leaves you with a choice of consuming \$80 today or, if you invest your profits, \$86.40 next year—reducing the reward to investing from 10 percent to 8 percent, just as in the personal saving example.

In contrast, under expensing the tax on business income has no effect on the return to investing. If you decide to spend all of your profit on the equipment, you face no tax liability today because you get to deduct the full purchase price of the equipment immediately, which makes your tax base today \$0. This leaves you with \$100 to invest in the equipment. Next year, you'll have \$110 of output to sell but no depreciation deductions. After paying a 20 percent tax on that \$110, you will have \$88 dollars left. Thus, you have a choice between consuming \$80 today or \$88 next year. This is, of course, a 10 percent return, the same as what you would get in the absence of taxes. Although our example could be made more complicated—for instance, by allowing the equipment to depreciate over several years—this would not change the basic conclusion that with expensing the rate of return to investment is the same as it would be in the absence of taxes.

Similarities and Differences between a Consumption Tax and a Wage Tax

Another way to eliminate the effect of taxes on the incentive to save and invest is to base tax liability on labor compensation, while exempting from tax all interest income and other returns from saving and investment. This approach is generally called a *wage tax*, although tax scholars sometimes refer to it as a *yield-exempt tax*, because the yield (i.e., return) from saving and investment does not trigger tax liability. Extending our example to this case is straightforward. Under a 20 percent wage tax, when you earn \$100, you receive \$80 after tax. Saving it at a 10 percent interest rate would give you \$88 next year, and there's no tax on the interest. So just like a consumption tax, the



wage tax gives you a choice of consuming \$80 today or \$88 in a year. Not only is there no impact on the incentive to save—the rate of return remains at 10 percent—but the amount you get to consume after tax is exactly the same in both periods under either the wage tax or the retail sales tax.

These two alternative methods for exempting the return to saving—a consumption tax and a wage tax—are reflected in the two different kinds of IRAs in the current system. A traditional IRA allows a deduction for contributions to the account, exempts capital income on the account from taxation while it accrues, and then taxes withdrawals (which are presumably made to finance consumption). A Roth IRA, on the other hand, does not allow an initial deduction, does exempt capital income on the assets in the account from tax, but does not impose tax on withdrawals. Thus, the traditional IRA works like a consumption tax does to eliminate the tax penalty on the return to saving, while the Roth IRA works like a wage tax. Except for some technicalities (for example, marginal tax rates may differ at the time of contribution and the time of withdrawal), the two approaches are similar in the same way that consumption and wage taxes are.

A crucial distinction between a wage tax and a consumption tax arises, however, in considering the transition from an income tax to one or the other. If we switch entirely to a consumption tax, any preexisting wealth that gets spent on consumption after the transition will be taxed as it is spent. However, if we eliminate the income tax and replace it with a wage tax, all preexisting wealth can be spent on consumption tax-free forever. This is indeed an enormous difference, as accumulated wealth in the United States was approximately \$55.6 trillion by the end of 2006.⁶

Not All “Capital Income” Represents the Return to Saving

So far, we’ve seen that a hallmark of both consumption taxes and wage taxes is that they exempt from taxation the return to saving. The return to saving can be thought of as the reward to deferring consumption. Certainly some of what we loosely call “capital income” (interest, dividends, capital gains, small-business income, and so on) represents returns to deferring consumption. But some of it represents something else entirely. In what follows, we consider what the other components of capital income represent and how they are treated under different tax systems.⁷



Consumption Taxes, Income Taxes, and Inflation

As we have noted before, some portion of capital income represents compensation for the fact that inflation erodes the real purchasing power of the underlying wealth. Consumption and wage taxes automatically exempt this compensation for inflation from tax. For instance, if half of the 10 percent interest in our examples above reflected inflation, the consumption and wage taxes wouldn't tax it because they exempt the full 10 percent from tax. In contrast, under our current income tax, to the extent that capital income is taxed, the inflationary component is taxed too. As we discuss in chapter 4, when inflation is more than a few percent a year, this can greatly increase the effective tax rate on the return to saving.

Compensation for inflation does not have to be taxed by an income tax. But designing an income tax that has this feature is difficult, while it is automatic under a consumption or wage tax. The key difference is that an income tax must distinguish what portion of capital income was compensation for inflation and what portion was not, instead of just exempting the whole thing. Accurately distinguishing the two would require not only a measure of the dollar amount of capital income, which is what we currently observe on the tax form, but also a measure of the value of the underlying wealth that generated the return, which can be administratively difficult to obtain, as well as the rate of inflation over the period the capital asset was held. For example, suppose that someone receives \$10 of interest income and the rate of inflation has been 5 percent. If that interest was earned on \$100 of wealth, we should be exempting 5 percent of \$100, or \$5. But if it were earned on \$125 of wealth, we should be exempting 5 percent of \$125, or \$6.25. If we don't know whether the underlying wealth was \$100 or \$125, we don't know exactly how much of the \$10 to exempt to offset the fact that, in addition to generating the \$10 of interest income, the real value of the investment has been eroded by the inflation. It would be possible to come up with rough approximations—for example, based on average interest rates in the economy in a particular year—but these would be just that, rough approximations. Obtaining a correct measure of real, inflation-adjusted, personal capital income would also require adjusting dividends and capital gains (i.e., the increase in value of assets held in one's portfolio) for inflation. At the business level, the value of depreciation deductions, interest deductions, and inventory allowances would also need to be adjusted for inflation. For similar



reasons as above, this is difficult in an income tax and unnecessary under a consumption or wage tax.

Consumption Taxes, Income Taxes, and Risk

Another important component of capital income is the compensation—or *premium*—for bearing risk. On average, riskier opportunities for saving and investment yield higher returns because people dislike risk and need to be compensated to be willing to accept it. For example, corporate stocks have historically earned a higher return, on average, than government bonds, partly because the return on stocks is much more variable. Therefore, there is a trade-off between risk and expected average return, and in choosing their portfolio savers need to choose some optimal balance between them. Some argue that, historically, the part of capital income that represents compensation for risk has been higher than the part that represents compensation for deferring consumption. But evidence also suggests that the return for deferring consumption has been increasing over time and the premium for bearing risk has been declining.⁸

Compensation for bearing risk is exempt from consumption taxes and wage taxes in exactly the same way as the compensation for inflation and for postponing consumption are. To illustrate, consider a risky asset that offers a 1-in-2 chance of earning a 40 percent return and a 1-in-2 chance of losing 10 percent of its value, so the return is expected to be 15 percent on average. Suppose once again that you've got \$100 before taxes available to save and there's a 20 percent consumption tax rate, so you have the option of consuming \$80 after taxes today. Under a consumption tax, if the money is saved, no tax is paid today, \$100 gets saved, and it yields either \$140 or \$90 next year. After paying the 20 percent consumption tax, you get to consume either \$112 or \$72 next year. Relative to the \$80 you can consume today, that still translates to a 40 percent return if you have good luck, a 10 percent loss if you have bad luck, and a 15 percent return on average. Under a wage tax, you pay tax immediately today even if you save, leaving \$80 to be saved. Next year you get to consume \$112 ($= 1.4 \times \80) if you have good luck or \$72 ($= 0.9 \times \80) if you have bad luck, exactly the same outcomes as under the consumption tax. So under either a consumption tax or a wage tax, your average return from investing in the risky asset is 15 percent, the same as it would be without any taxes, and there is no distortion to your decisions regarding how much risk to take. Comparing a consumption tax to a wage tax, the people who are lucky after



the fact remit more tax, and the unlucky people remit less tax. But this distinction is in fact meaningless. What people actually get to consume either today or tomorrow, whether they have good luck or bad, is exactly the same under both taxes.

A more surprising principle is that under a proportional income tax that allows full deductions for losses, the compensation for risk is also effectively exempt from tax. The argument for why, which dates back to an article written by economists Evsey Domar and Richard Musgrave in the early 1940s, is based on the key insight that the income tax absorbs some of the risk from your portfolio. If an investment turns out well, the income tax takes away some of the resulting capital income, but if it turns out badly, the tax system partly compensates you by allowing you to deduct the loss, thus reducing tax liability. Thus, the income tax reduces both the expected after-tax return *and* the variance in the after-tax return. Domar and Musgrave showed that the reduction in the expected after-tax return is the fair market price for the reduction in risk, so that even though on average the government collects revenue from imposing income tax on the portion of the return that represents compensation for bearing risk, that part of the tax does not place a burden on the taxpayer. In fact, by increasing the share of one's portfolio invested in risky assets, an investor can exactly replicate the risk-return possibilities that would have been available if only the risk-free portion of returns were taxed.⁹ Since the income tax on the risky portion of the return leaves you with exactly the same opportunities as you would have had with no tax at all on that portion of the return, it is in a sense equivalent to not taxing that portion of the return at all.

This example relies on a presumption that the income tax treats gains and losses symmetrically, which in reality it does not. Progressive tax rates and limitations on deductibility of losses both mean that the government takes a larger share of capital income if you have good luck with your portfolio than if you have bad luck. However, this is not an inherent or unique feature of income taxes. The reasons for these features (the desire for ex post vertical equity and the need to limit tax avoidance schemes) can apply to consumption tax plans as well.¹⁰ A consumption tax with these features would also tax some of the compensation for bearing risk.

“Supernormal” Returns and Mislabeled Labor Compensation

Aside from compensation for postponement of consumption, inflation, and risk, some of what we loosely call “capital income” may reflect



a variety of other things as well, such as returns to innovation, returns to entrepreneurial skill and effort, returns to establishing a monopoly in some market, returns to early units of investment that may be more productive than the last “marginal” units of investment, or simply mislabeled returns to labor effort. What all of these have in common from an economic point of view is that they are *supernormal returns*—meaning they produce income that is larger than the “normal” return that would be earned on a marginal investment in capital, which is competed down to a fairly low level. It turns out that these supernormal returns are effectively taxed by both income taxes and consumption taxes. A wage tax, on the other hand, would not tax them.

A highly visible example of a supernormal return would be Bill Gates’s income, which comes mainly from his share of the profits of Microsoft. Only a tiny portion of that income represents a reward for postponing consumption or taking on risk. Mostly, it represents some combination of returns to Bill Gates’s labor (including effort put into innovation and entrepreneurship) and returns arising from being in the right place at the right time and gaining an early lead in what turned out to be a fantastically lucrative market that naturally tends toward monopoly (because computer software tends to be more valuable to a user when most other people are using the same software). Under the current income tax, these returns are indeed taxed, mainly by the corporate income tax and also by the personal income tax to the extent that Gates realizes any of his capital gains by selling shares or receives dividends from Microsoft.

A wage tax, on the other hand, would never tax these returns as long as Microsoft continued to pay Gates only a nominal salary. Under a wage tax, relabeling your labor compensation as capital income would be an easy avenue to escape taxation altogether. The difficulty of distinguishing what is labor income and what is capital income is an important reason that a pure wage tax would end up being highly inequitable and costly to enforce.¹¹

In contrast, a consumption tax does indeed tax supernormal returns and mislabeled returns to labor (like those received by Bill Gates) in the same way as an income tax does. One way to look at this is that no matter what these returns are labeled, they will eventually trigger tax liability when the recipient, or his or her heirs, eventually consume more because of these returns. For consumption taxes like a flat tax that



have a business tax component, the normal return to saving and investment is exempt because investments in physical capital (such as productive machinery and factories) are expensed (deducted immediately). The factors that produce supernormal returns for someone like Bill Gates lead to big increases in the taxable revenues of his company but do nothing to increase the size of the expensing deduction that the company gets to take early on: the supernormal returns did not result from larger than usual investments in physical capital. Similarly, under the Hall-Rabushka flat tax, if a firm mislabels the labor compensation of its proprietor as “dividends” or “capital gains,” no tax will be paid on that income at the personal level, but tax is paid at the business level, and the mislabeling produces no expensing deduction for capital investment to offset it. Thus, as long as the tax rates on the personal and business bases in the flat tax are the same, mislabeling labor income as returns to capital and vice versa make no difference to the total tax liability.

How Different Is the Current Income Tax from a Consumption Tax?

The bottom line of the above discussion is that an income tax taxes the return to saving while a consumption tax does not. Compensation for the loss in real value of assets caused by inflation need not be taxed under either type of tax, but it is administratively easier to exempt it from a consumption tax. Compensation for bearing risk need not be taxed under either approach: the features that make it partly taxable (including progressive rates and limitations on loss deductibility) might apply under either type of tax. Finally, “supernormal” returns and labor compensation mislabeled as capital income are generally taxable under either approach. Thus, much of what is generally thought of as capital income can be treated similarly by income taxes and consumption taxes. In sharp contrast, a wage tax, would completely exempt *all* of these forms of capital income.

The difference between a consumption tax and the system that we have now is mitigated further by the fact that, as we detail in chapter 2, our current tax system is in many ways a hybrid between an income tax and a consumption tax. For example, although revenue is collected from taxing nominal (not real, inflation-adjusted) interest income, even more revenue may be lost due to the *deduction* of nominal interest payments, both because those deducting interest payments tend to be in higher tax brackets than those receiving interest and because so much interest



income is sheltered from tax (perhaps, because it is earned in tax-exempt pension accounts or IRAs). Large portions of dividend and capital gains income are excluded from personal tax, and generous depreciation allowances further reduce the effective tax on capital income.

One rough metric of how different the current income tax is from a consumption tax is the amount of revenue that would be lost if the tax base was switched from income to consumption, while leaving the other features of the current system intact. To see this, consider that the existing income tax could be almost entirely converted to a type of consumption tax in just three steps: (1) exempt all interest, dividends, and capital gains from tax; (2) replace depreciation deductions with expensing of all new investment, while eliminating depreciation deductions for past investment; and (3) eliminate all interest deductions. One recent study calculated that taking these steps in 2004, while leaving the rest of the tax code (including the rate structure) the way it is, would have reduced revenues from personal and corporate income taxes by a total of \$64 billion that year, or approximately 6 percent of income tax revenues.¹² Thus, in 2004, taken altogether, the parts of our tax system that make it a hybrid between an income tax and a consumption tax, rather than a pure consumption tax, raised only a modest amount of net revenue. In that sense, we are not that far from a consumption tax, although our system is still rife with many distortions to incentives that would be avoided under a clean consumption tax. This should not be taken to mean that transforming our income tax into a consumption tax would be easy, though. The calculation above relies heavily on the politically challenging steps of eliminating depreciation deductions and/or interest deductions. Skipping either or both of those steps would make the switch to consumption taxation very costly indeed in terms of revenue.

Now that we've examined how income taxes and consumption taxes differ in how they treat capital income, let's consider how they treat the other major source of income, labor.

Consumption Taxes, Income Taxes, and the Incentive to Work

One thing that consumption taxes and income taxes share is that they reduce the reward for working (the quantity of goods and services that can be purchased per hour of work). The reward for working can be reduced in either of two ways—by lowering take-home pay and holding fixed the prices of the things you buy or by increasing the prices of



everything you buy and keeping take-home pay the same. A comprehensive 20 percent tax on wages and salaries has about the same effect on reducing the reward for working as a comprehensive 20 percent sales tax.¹³

It is sometimes argued that a consumption tax reduces the incentive to work *more* than an income tax does. The argument goes like this. Total consumption is smaller than the total amount of income. Therefore, the tax rates required to raise a given amount of revenue must be higher under a consumption tax than under an income tax with an equally broad base, and a higher tax rate means a greater disincentive to work. The problem with this argument is that it ignores the fact that, because a consumption tax allows you to earn higher after-tax returns on the portion of your labor income that you save, working becomes more financially attractive.¹⁴ This helps offset the effect of the higher tax rates (on an apparently smaller base) that would be required under a consumption tax.

Is a Consumption Tax More or Less Fair Than an Income Tax?

On a year-by-year basis, a consumption-based tax can appear to be much more regressive than an income tax with the same rate structure. This is because people who have low incomes in any given year on average spend a very high fraction of their incomes—often more than their income—while people with high incomes in a given year save a relatively larger portion of theirs.¹⁵ The temptation might be to conclude that shifting from an income tax to a consumption tax with the same rate structure would greatly increase the tax burden on low-income people relative to high-income people.

This argument is somewhat misleading, though, because looking at a snapshot of income-saving patterns for a single year of data significantly overstates the regressivity of the consumption tax compared to an income tax. First of all, in any given year some people with temporarily low income continue spending at a rate corresponding to their usual income. Conversely, some people who had an exceptionally good income-earning year might not expect their good fortune to continue and therefore keep their spending well below their income, saving up for the years of relatively bad fortune. In both cases, one year's income is not a good measure of how well off those people are, and the fact that the consumption-income ratio varies widely across persons, being high for the temporarily low-income people and low for the



temporarily high-income people, overstates how much more regressive a consumption tax is relative to an income tax.

Second, most people have a natural life-cycle pattern of earning and saving. In the early years of working life, family expenses are pressing and income relatively low, so savings are minimal or even negative, as families borrow to finance consumption. In the later working years, incomes have grown to the point where many families begin to save for their retirement and higher education for their children. In retirement, the pattern reverses again, as people live off their accumulated savings, spending more than they earn. Looking across people of different ages, low-income people (on average the young and old) would appear to do little (or none or even negative) saving, and high-income people would seem to be relatively big savers. That picture would be misleading because—leaving aside for a moment bequests, inheritances, other gifts, and government transfers—over a lifetime people cannot spend more than they earn, and over a lifetime people end up spending all of what they earn.¹⁶ Thus, from a lifetime perspective it doesn't make a huge difference whether the tax base is income or consumption: it all adds up the same. Big savers age 45 to 65, who might appear to be getting off easy under a consumption tax because they consume relatively little of their annual income, will eventually pay more tax when they take their trip around the world and otherwise live the high life when they retire. A single-rate tax on consumption and a single-rate tax on labor income both end up being single-rate levies on lifetime resources. Similarly, over a lifetime, a consumption tax with graduated rates could in principle achieve about the same degree of progressivity as a graduated income tax: one is not inherently more progressive than the other. Rather, the degree of progressivity depends largely on the kind of tax rates that we impose, which is in principle a separate issue than the tax base the rates apply to.

Many economists would also argue that, again ignoring bequests and inheritances, a consumption tax is more horizontally equitable than a comprehensive income tax. This is because when comparing families with equal lifetime incomes, a consumption tax levies the same present value of taxation regardless of a family's saving propensity. In contrast, under an income tax savers pay more tax than nonsavers because the return to their saving triggers an additional tax liability that does not arise under a consumption tax.

One caveat to this argument is that a consumption tax that does not include bequests or inheritances in the tax base will lower the average



tax rate over a lifetime on those families that pass on wealth to their heirs. Evidence suggests that people who have higher lifetime incomes leave a larger share of their lifetime resources as bequests than do other people. This turns a flat-rate consumption tax into a somewhat regressive tax on lifetime resources. For this reason, some economists who favor a consumption tax do so only if a bequest is treated (that is, taxed) as if it were an act of consumption by the bequeathor. A tax on bequests, however, might be viewed as a double tax on deferring consumption, as a tax is levied both when the bequest is made and then again when the heirs consume it. If people leave bequests because they value the consumption of their heirs, then this would reduce the incentive to save for bequests in much the same way as an income tax does. We consider these issues in more detail when we discuss the estate tax in chapter 8.¹⁷

Fairness and the Transition to a Consumption Tax

Whenever the tax system changes, some people are bound to lose in the transition, and others are bound to gain. The bigger the tax change, the bigger the likely size of the windfall gains and losses. Replacing the current tax system with any kind of consumption tax is surely a massive change, and therefore potentially large windfall gains and losses would occur in the transition.

The biggest issue in the transition to a consumption tax is the treatment of preexisting wealth. The transition to a new tax system could mean either a major windfall or a substantial loss for the owners of wealth, depending on how the change is implemented and in what form the wealth is held.

To see this, recall that either a wage tax or a consumption tax removes the impact of taxation on the incentive to save. If we were to replace our income tax with a wage tax, not only could wealth owners earn interest, dividends, and capital gains tax-free, but their wealth could (as under the income tax) be spent tax-free as well. This change would clearly provide a big gain to holders of existing wealth, who otherwise would have had to pay tax on the returns to that wealth. It is a *windfall* gain because these people gain financially from the new tax rules, even if they don't do any extra new saving at all.

On the other hand, adopting a pure consumption tax could impose a large burden on owners of existing wealth. Suppose that we adopted a 20 percent consumption tax and it caused the prices of all goods and services to rise by 20 percent. The purchasing power of labor income



would decline by 20 percent relative to a situation with no taxes, just as it would under a wage tax of the same rate. The difference is that, upon adopting a consumption tax, the purchasing power of not only labor income—but also of *all* existing wealth—falls by 20 percent. For example, suppose you had built up a substantial savings account before the imposition of the new consumption tax. When you withdraw money from the account and spend it, it would be taxed at the full 20 percent rate by a consumption tax but not taxed at all by a wage tax.

Setting aside for a moment temporary inflexibilities in, for example, wage contracts (we address these later), whether the overall level of prices changes or not does not materially affect this story.¹⁸ Even if prices do not rise at all, moving to a consumption tax would cause the purchasing power of both wages and existing wealth to decline by an average of 20 percent relative to a situation with no taxes. Nominal wages would be forced down because firms would be earning 20 percent less, after taxes, from the output produced by workers. The nominal value of existing capital assets—in the form of, for example, share prices—which constitute much of old wealth, would also decline because the output they produce provides 20 percent less in after-tax revenues.

The effect of eliminating the current income tax needs also to be considered. For wages, the offset would be roughly one-for-one: the elimination of income taxes on wages would compensate for the new burden of consumption taxes on wage earners. Different wage earners might gain or lose depending on how the progressivity of the tax burden on wages changes, but this is a separate issue.

The future benefits from eliminating income taxes on capital, however, are unlikely to cancel out the immediate one-time tax on existing wealth from instituting a consumption tax. Eliminating double taxation of all future dividends and capital gains would provide a counterbalancing gain to corporate shares, as would forgiving the tax on all previously accumulated capital gains.¹⁹ But eliminating the income tax also imposes a loss on owners of capital assets because firms would no longer be able to take depreciation deductions on previously purchased capital assets to offset taxes on the output they produce. All new investment goods, by contrast, could be purchased tax-free or deducted from the business tax base immediately, which would offset the taxes owed on the future returns to those new assets. As a result, the demand for existing assets would fall sharply relative to the demand for new assets, reducing the value of existing wealth. The net impact



on the value of existing wealth from a switch to a consumption tax depends on many factors and is hard to predict.²⁰ All things considered, however, the most likely scenario is that owners of existing wealth would take a one-time hit in the switch from the income tax to a consumption tax.

The potential windfall tax on existing wealth can be relieved to some extent by building transitional provisions into the reform. However, the more transition relief that is provided to existing assets in the switch to any consumption tax, the more it becomes like a wage tax. For example, allowing businesses to continue taking depreciation deductions for past investments would greatly reduce the revenue raised from holders of preexisting wealth. Taxes on labor compensation would have to pick up the slack. A similar result occurs in the transition to a personal consumption tax if individuals are allowed to deduct some or all of their existing wealth in the same way as they could deduct new saving. To the extent they can, their tax base is less than their consumption done after the imposition of the tax. A similar outcome results from incremental changes in the income tax that exempt more and more of the return to saving from taxation without imposing any transitional burden on preexisting wealth. In sum, whether moving toward consumption taxation provides a windfall gain or loss to existing wealth holders depends crucially on the transition provisions that are adopted.

What characterizes the owners of wealth who will be affected most by these transitional gains or losses? They are, on average, very wealthy; for example, the richest 1 percent of the population owns approximately one-third of the nation's private wealth.²¹ They are also disproportionately older people; in 2004, people age 55 or over were about 31 percent of the adult population, but households headed by someone 55 or over owned 57 percent of the nation's wealth.²² If Social Security benefits are indexed to compensate for any price change, as they probably would be, then the impact is concentrated on the better-off older people.

Whether it is "fair" to hit wealth holders, elderly or not, with such an unexpected burden is an equity issue about which economic analysis cannot be decisive. The fact that moving to a consumption tax with no transition relief places a burden on the wealthy is probably appealing to many people. But experience suggests that wealth owners and their representatives—such as corporations—would be sure to push very hard for transition relief, and the political system is likely to



accommodate them. Indeed, an article in *Fortune* magazine predicts that “Were Washington to disallow deductions [for depreciation on existing capital], every CEO-laden corporate jet in America would commence strafing Capitol Hill.”²³ The other major group hit especially hard by the transitional burden—retirees with significant wealth outside of Social Security—would also likely see their interests defended by strong lobbying, a powerful voting bloc, and general public sympathy, so some kind of transitional relief could be expected from this quarter as well.

Because many important prices are inflexible in the short run, the transition to a consumption tax can also lead to a variety of other, often capricious redistributions. Depending on whether the Federal Reserve allows the price level to change, all sorts of redistributions can occur surrounding contracts written in fixed nominal terms, such as bonds, mortgages, and long-term wage contracts. For example, if a switch to a consumption tax is accompanied by an increase in consumer prices, bondholders suffer the same one-time loss as stockholders because the real value of their bonds declines. On the other hand, if prices do not rise, the switch to a consumption tax will not hurt bondholders because firms are legally obligated to pay them the full nominal value of their bonds. In this case, the owners of businesses who borrowed from those bondholders are hit hard because they cannot pass through the cost of lost depreciation deductions to the bondholders. Some of these firms might be forced into bankruptcy or layoffs, hurting the employees as well.

One other notable issue of transitional fairness concerns government transfer payments that are not automatically indexed for price changes, such as food stamps and welfare. If a consumption tax is accompanied by a price increase, and transfer payments are not increased in value to keep real benefits constant, the tax change levies an extra burden on the poor.

Economic Reasons for Switching to a Consumption Tax

What would be the long-term economic impact of switching to a consumption tax? The most talked-about potential benefits arise from increasing the incentive to save and invest. A less publicized but perhaps more important benefit is that a consumption tax could make it easier to “level the playing field” among different types of investment. Finally, the windfall of tax burdens that arise in the transition to



a new system could have important economic effects, apart from whether they are fair or not.

Switching to a consumption tax would remove any negative impact of the tax system on the incentive to save and invest. Thus, a consumption tax distorts one less decision—and a critical one, to boot—compared to an income tax. By itself, this is a step toward a more efficient tax system because individuals' choices between current and future consumption will more closely reflect the social return to saving rather than the influence of the tax system. However, it is not necessarily true that this makes a consumption tax more economically efficient overall than an income tax; the tax that alters the smallest number of relative prices is not necessarily the most efficient.

In theory, a reform that removes the tax distortion to saving decisions while raising the same revenue should increase saving and capital accumulation, which many economists argue would be beneficial. By providing workers better machines and tools to work with, increased capital accumulation could improve productivity and long-run living standards. As shown in chapter 4, however, the magnitude of any increase in saving that would arise from improved incentives is highly uncertain. This is a crucially important question on which the available evidence is simply not very good. In recent history, the periods with the highest incentives to save coincided paradoxically with the *lowest* rates of saving. This might have occurred for many different reasons, each with different policy implications, and it is difficult to distinguish empirically among these reasons. Nonetheless, the best guess of most economists is that private saving is probably not very responsive to the after-tax rate of return, so that switching to a consumption tax would be unlikely to increase the quantity of saving much. Because we can find more direct ways to increase national saving (for example, by reducing the budget deficit), the likely but not guaranteed prospect of a somewhat higher saving rate does not appear to be, by itself, a reason to undertake a wholesale transformation of the tax system.

A second possible benefit of a consumption tax is that achieving “neutrality” or “uniformity” in the tax treatment of various types of investment would be made easier. Under the current income tax, different types of investment are effectively taxed at varying rates. For example, corporate business investments are taxed more heavily than noncorporate business investments, which are in turn taxed more heavily than investments in owner-occupied housing. Investments in



certain types of capital equipment or in certain lines of business are capriciously favored relative to others by depreciation schedules that are more front-loaded, or accelerated, than true economic depreciation. Business endeavors that can be packaged into assets that appreciate in value are more attractive than those that pay their returns in the form of dividends or interest because of the preferential treatment of capital gains. Inflation can introduce arbitrary variation in tax rates. All of these deviations from “neutrality” are economically harmful because they attract investment into the tax-favored investments even though their social return is not higher than the social return to the investments that are not tax-favored.

A consumption tax would eliminate all of these distortions to the choice among different types of investment because it would equalize the effective marginal tax rate on the returns to all investment to a single uniform rate—zero. In principle, an income tax could also be made more neutral, taxing the returns to investment at a uniform but positive rate. Thus, to some extent, this is a separate issue from the choice between an income tax and a consumption tax: it’s a matter of having a “clean,” or uniform, base. But this particular aspect of uniformity is much easier to achieve in a consumption tax than in an income tax. Some of the existing distortions are hard to avoid in an income tax because capital income is difficult to measure accurately. For example, measuring economic depreciation exactly right is impossible, so some distortions will inevitably be caused by favoring investments in assets whose depreciation for tax purposes exceeds true depreciation. Similarly, including all capital gains in taxable income as they accrue is probably infeasible, so appreciating assets will be favored. It is theoretically possible to adjust the measurement of capital income for inflation, but in practice doing so would be complicated and inevitably imperfect. Because a consumption tax eliminates the need to address any of these issues, it can achieve this aspect of uniformity in a simple way.

The transition to a consumption tax can also have significant efficiency consequences. For example, from an efficiency point of view, placing a surprise, one-time tax on holders of existing wealth, as moving to a consumption tax could, is not unattractive. Raising revenue from the returns to *past* investments has no effect on the incentive to invest, to work, or to do anything else. Shifting tax burdens onto the elderly—whose work and saving decisions are nearly finished—also avoids costly disincentives. Moreover, because the elderly have a relatively



low propensity to save, transferring some of the tax burden onto them and away from others could give a boost to national saving. This last effect is likely to be offset, however, by the fact that the wealthy as a whole have a much higher propensity to save than everyone else, and placing an extra burden on them through a consumption tax might reduce their contribution to national saving.²⁴ As we see in the next chapter, models that attempt to work out the economic impact of switching to a consumption tax suggest that, all things considered, a very large portion of any economic gain from the switch comes from replacing other, distorting, taxes with this one-time tax on already accumulated wealth.

Any potential economic benefit from the transitional tax on preexisting wealth that would accompany a switch to a consumption tax depends on whether people will really believe that this is a “one-time only” strategy. If a surprise one-time tax on wealth is so economically efficient if we do it “just this one time,” politicians might well be tempted to try it again some time in the future. If doing it once increases people’s estimates of the probability that it will happen again, then the “surprise” tax on wealth creates a disincentive to accumulate wealth.

As we have emphasized, we now have a messy hybrid between an income tax and a consumption tax. This is a mixed blessing for the case for switching entirely to a consumption tax. To the extent that saving already receives consumption tax treatment due to pensions, IRAs, and the like, the potential increase in saving and corresponding economic benefits that could arise from going all the way to a consumption tax are apparently reduced. This should not be overstated, however. Although a large share of capital income is already sheltered from taxation, a large share of the capital income that would arise from *additional* saving, which matters for decisions about how much to save, is still taxed. This is largely because pensions, IRAs, and similar plans shelter saving only up to some limit. For anyone whose saving is at or above those limits, the marginal incentive to save is still distorted by tax. On the positive side, this means that switching entirely to a consumption tax would still improve incentives substantially for the people who do most of the saving—and at a lower cost in terms of lost revenue (which must be recouped somehow) than if we were starting from a pure income tax. But it also means that the capital income that remains unsheltered from the current tax disproportionately goes to high-income people, which heavily influences the distributional impact of such a reform.



Simplification and Enforcement Aspects of a Consumption Tax

Even with all the compromises that we make in our current system, calculating and reporting capital income can be a complex and burdensome process. These compromises, in turn, create opportunities to achieve tax savings in complicated and socially unproductive ways, such as devising schemes to rearrange financial and business transactions. In principle, measuring consumption accurately is simpler than measuring even a compromised version of income, mainly because the need for measuring capital income can be completely avoided. But in practice, a consumption tax could end up being just as, or even more, difficult to administer and comply with than the current income tax. The simplicity and enforceability of moving to a consumption tax depend crucially on which approach is chosen and how it is operated. These issues are addressed fully in chapter 7.

Here, too, transitional issues are critically important. Moving cold turkey to a consumption tax like a VAT or flat tax could immediately make the taxpaying process simpler. But the pleadings of those likely to suffer windfall losses will be difficult to resist, so a switch to a consumption tax is likely to offer various forms of transition relief. Special transitional provisions can be exceedingly complex because they require the simultaneous operation of parallel tax systems, at least for a while. For example, interest on mortgages taken out before the switch might remain deductible, while interest on new mortgages is not. The real choice is not between the current system and a clean consumption tax but between the current system and a new tax encumbered by as yet unspecified rules for how to get from here to there, which can be unaesthetic at best and complicated and loophole-ridden at worst.

One other transitional issue has important implications for evaluating a new system. Depending on how it is implemented, the switch to a consumption tax can lead to enormous incentives to postpone or speed up transactions around the date of switchover. Firms will postpone investment until the expensing rules are in place, and under some plans, individuals will rush to consume as much as possible before the implementation date. This could cost a great deal of revenue and raises a very delicate problem of minimizing the short-run disruption at the time of transition.

A Clean Tax Base

A third element of proposals for fundamental tax reform is to eliminate many or all of the features that provide special treatment of particular



types of consumption or investment. The aim is variously described as “flattening” or “broadening” the tax base, “leveling the playing field,” or making the tax system more “neutral.” We refer to this aspect of flatness as a *clean* tax base. (At least the connotation of *clean* is less ambiguous than that of *flat*.)

One important source of messiness in the tax base has to do with compromises or peculiarities about the taxation of capital income, such as the failure to index capital income for inflation and the double taxation of corporate income. These can be important sources of inefficiency, inequity, and complexity. As discussed above, a consumption tax could sweep away many of these problems.

A second source of messiness in the tax base concerns preferences for particular types of expenditures or activities that are deemed worthy of special treatment. Important examples are housing, health care, charity, taxation and borrowing by state and local governments, and education. Although we concentrate here on a few important examples, the current personal and corporate income taxes feature scores of other deviations from a clean-base system, from incentives to invest in low-income housing to tax breaks for the production of ethanol.²⁵ Each needs to be evaluated on its own merits.

Some ambitious proposals for overhauling the tax system, such as the Hall-Rabushka flat tax, would eliminate *all* preferences of this sort. Other proposals eliminate some but not all of these features. In contrast, while recent U.S. policy changes are moving the tax system closer to a consumption base and flatter rates, at the same time they are rapidly *increasing* the extent and generosity of deductions, exclusions, and credits. Any attempt to remove the most cherished and politically entrenched of these tax goodies would undoubtedly engender a fierce political fight, as the negative reaction to the 2005 President’s Advisory Panel on Federal Tax Reform showed. Were we to move to a consumption tax that features a personal tax return, such as the flat tax, it could easily end up retaining many current preferences or even adding new ones. And even though consumption taxes that have no personal tax returns, such as a retail sales tax or VAT, could no longer contain the preferences in their current form, similar ones could be reintroduced by exempting purchases of preferred items.

Throughout this book, we have discussed a host of reasons why special preferences in the tax base can be a problem. Every preference is a penalty for someone else because it requires tax rates to be higher than otherwise. Tax deductions, in particular are a regressive way to subsidize activities: people with larger incomes receive bigger



subsidies from deductions because on average they engage more in the deductible activities and have higher marginal tax rates. Those who don't itemize get no benefit at all from the deductions. Except in special cases, tax preferences are inefficient because they create an incentive to engage "too much" in the lightly taxed activity and too little in other activities, relative to what the free market would dictate. And finally, they are a big reason that our tax code is so complicated.

For all of these reasons, the burden of proof should rest on those who defend deviations from a clean tax base. There are two main acceptable lines of defense. First, allowing certain deductions can make the tax system fairer if by so doing the tax base provides a more accurate measurement of well-being. Second, a tax preference can improve economic efficiency if it corrects a significant market failure—that is, a situation where the incentives in the market do not lead to an economically efficient outcome by themselves. An example would be an externality, which occurs when an activity generates important benefits or costs for others that are not reflected in the incentives faced by the individual undertaking the action. Even if a significant market failure is identified, the tax preference must be the best alternative for dealing with it.²⁶ As specific exceptions to a clean base are examined, each of them is evaluated in light of these criteria.

Even if some preferences meet these criteria, one powerful argument for maintaining a clean base is that the political system is incapable of distinguishing legitimate arguments from illegitimate ones and often succumbs to the political clout of powerful pleaders. Once any preference is allowed, we may begin to slide down the slippery slope to more preferences.

Housing and the Mortgage Interest Deduction

Owner-occupied housing is favored by our tax system in a number of ways. By far, the deduction for home mortgage interest payments is the most visible and traditionally the most politically sacred aspect of these. Notably, its "sacredness" is now being challenged, for it is swept away in some of the more ambitious tax reform proposals. The 2005 presidential tax commission recommended major modifications.

The home mortgage interest deduction is an expensive one, costing \$80 billion in revenues in fiscal year 2007. It requires personal tax rates to be about 6.8 percent (not 6.8 percentage *points*) higher than they otherwise could be.²⁷ So a lot is at stake in this debate. What are the issues?



First of all, let's consider whether there is any good economic reason to favor housing over other types of consumption or investment. To answer yes to this question requires demonstrating that an owner-occupied house provides important benefits to people other than the residents themselves; that the residents themselves take pleasure in ownership is an inadequate argument. Although undoubtedly neighbors prefer to gaze out their window at a well-kept rather than ramshackle house and owner-occupiers arguably maintain their houses better than the combination of renters and landlords, these benefits are certainly quite localized and probably fairly small. Sometimes vague appeals are made to the role of homeownership in maintaining a strong democracy, but these arguments are not convincing. Remember that to the extent that the tax system attracts investment into housing, it diverts funds from other business investments, which lowers the productivity of workers and the wages that businesses can profitably pay them. Why isn't broad stock ownership a healthy aspect of democracy? Preferential treatment of housing, in general, is difficult to justify on economic grounds.²⁸

Some evidence shows that homeowners are more likely to engage in home maintenance, especially gardening. Homeowners are also more apt to join organizations and socialize with their neighbors compared to renters. In addition, homeowners are more likely to be politically informed and active than renters. All of these activities may create positive externalities. The clear correlation between homeownership and these behaviors may not be indicative of a causal relationship, however, and could instead arise from personal characteristics that make certain people more likely to be homeowners and also make them join more clubs, be more conscientious gardeners, and so on.²⁹

In the specific case of the home mortgage interest deduction, there is a second possible rationale: it is needed to correctly measure the tax base in a comprehensive income tax. After all, to achieve a comprehensive measure of income, interest payments ought to be deductible, just as interest receipts are taxable. The catch to this argument is that, under a comprehensive income tax, the rental value of owner-occupied housing, net of depreciation and maintenance expenses, should also be subject to tax. To be sure, most homeowners don't think of the rental value of their home as income in the same way that they think about, say, their salary. But the failure to include the services provided by housing in the tax base, in conjunction with the deductibility of mortgage interest, adds up to a big preference for residential housing.



Consider a family that is trying to decide whether to buy a \$300,000 house or instead to buy a more modest \$200,000 house and invest the extra \$100,000 in the stock market. Suppose further that the annual rent for such houses would be 10 percent of their value and that stocks provide an annual pretax return of 10 percent, so that each would be an equally attractive investment in the absence of taxes. Living in the more expensive house will certainly make the family better off; they'll have nicer living quarters, more rooms, a better view, and so on. How much better off they are per year is approximated by the rental value of the extra housing, or \$10,000 per year. Investing the \$100,000 in stocks, on the other hand, will yield considerably less than \$10,000 per year because of the taxes that would be due on the investment return.

The preferential tax treatment could tip the scales in favor of investing in the more expensive house. As a result, in some cases nonhousing investments are passed up in favor of more expensive homes even though, taxes aside, the return to these investments equals or exceeds the value of the housing services; from a social point of view, this is wasteful. Allowing the deductibility of mortgage interest exacerbates this problem because it enables homeowners to use tax-deductible debt to finance an investment for which the return, the rental value of the housing, is untaxed. In addition, families and individuals who, for one reason or another, prefer to rent housing, end up being penalized. Because the net rental income of landlords is taxed, rental housing does not get the same preferential tax treatment afforded to owner-occupied housing. This generates an additional source of inefficiency. Not only is there an excessive amount of housing, but some households are induced to own housing when, taxes aside, they would find it more attractive to rent housing.

Even if we wanted to, however, there is no clean and easy way to put owner-occupied housing on a level playing field with other investments because taxing the net rental value of housing would be complicated and imprecise. Moreover, it would undoubtedly be resisted strenuously by the public, so it is probably not a practical option. It is not inconceivable, however, as several European countries have attempted to do this, albeit in a very rough fashion.³⁰

A simpler approach to reducing the tax preference for owner-occupied housing (that is, the penalty for everything else) would be to eliminate the deduction for home mortgage interest; Canada and Germany, and the United Kingdom are three countries that take this



approach. This would certainly make mortgage-financed housing less attractive, reducing the inefficient bias toward housing investment. But this is not a flawless option because it would eliminate the bias only for taxpayers who must borrow to finance their houses. For those who are wealthy enough (or who have wealthy-enough relatives) to pay cash, part of the cost of housing would still be the forgone, taxed, return on an alternative investment. The taxation of the alternative investment still makes buying one's own house look more attractive than otherwise. Thus, without a mortgage deduction, it would cost mortgage holders the pretax rate of interest to own housing (because interest is not deductible) but for wealthy individuals who need not borrow the cost would be the lower after-tax interest rate.

No such dilemma arises under a consumption tax. Although under some variants of the Hall-Rabushka flat tax proposed in Congress the deduction for mortgage interest is retained, it is completely incongruous in such a consumption tax because other interest payments are not deductible and all forms of capital income are untaxed.³¹ A consumption tax that treats purchasing housing just like other purchases and does not allow mortgage interest deductions would be a clean and simple way of eliminating the existing bias toward housing.

Tax preferences for owner-occupied housing are often defended on the grounds that their elimination would lead to unfair transitional effects. In particular, homeowners are understandably concerned about what eliminating the mortgage interest deduction would do to house prices. By itself, eliminating the home mortgage deduction makes it more expensive to buy a house for anyone who relies on mortgage financing and could expect to itemize deductions. This would tend to reduce the demand for housing, which in turn would cause housing prices to drop. The greatest decline in demand would be for the high-priced homes that itemizers in high tax brackets generally desire, so that eliminating the tax preference for housing should shift demand from more to less expensive homes.³²

The historical experience in the United States and other countries casts doubt on predictions of large impacts on either housing prices or the extent of homeownership. After all, dire predictions for the housing market were made about the Tax Reform Act of 1986, which made the deductibility of mortgage interest less valuable by lowering marginal tax rates, but it has been difficult to discern any negative effects on housing prices from that reform.³³ Canada has no deduction for mortgage interest at all but has almost the same homeownership rate as the



United States.³⁴ William Gale of the Brookings Institution points out that mortgage interest subsidies have been reduced dramatically in the United Kingdom since the mid-1970s, yet homeownership rates, mortgage debt, and housing as a share of the capital stock actually grew faster than in the United States.³⁵ Gale argues that a fundamental reform that eliminates all tax preferences for housing and moves to a consumption tax could, though, have a fairly large impact, with real house prices falling by about 7 to 10 percent in the short run and from 2 to 6 percent in the long run.

Health Care and Health Insurance

Our tax code provides very favorable treatment to health care and health insurance, and any attempt to change this in a reform effort will undoubtedly be a major source of contention. The most important preference in this area is the exclusion from tax of employer contributions to employee medical insurance plans. In 2007, this treatment cost an estimated \$141 billion in revenues, requiring personal tax rates to be 12.1 percent (not percentage points) higher than otherwise. In addition, an itemized deduction for large out-of-pocket medical care expenditures (those that exceed 7.5 percent of adjusted gross income) costs about \$4 billion in revenues.³⁶

Does health care merit special treatment? To be sure, health insurance is plagued by serious market failures arising from imperfect information.³⁷ First, a problem of *adverse selection* arises because individuals know more about their health risks and status than insurers do. Insurance companies charge premiums based on the average level of risk for a particular population and in some cases a perfunctory medical examination. Some healthier, lower-risk people will find this price unattractive and leave the market, which increases the average risk level of people left in the insurance pool. This in turn pushes up premiums, driving out even more low-risk people, and a vicious cycle ensues. As a result, many people end up without insurance, even though the insurance could be both valuable to the buyer and profitable for the insurance company in the absence of the adverse selection problem.

The system of employer-provided insurance helps mitigate the adverse selection problem, since employees work for firms for reasons generally unrelated to health risks, so that an insurance provider has reason to expect that any firm's employees are a mix of high-risk and low-risk people. Nevertheless, adverse selection is a serious problem



for people who are not covered by employer-provided insurance and are charged high premium rates. It is also arguably responsible for the almost complete absence of a market for long-term contracts for health insurance that would be fully portable across jobs and would not skyrocket in price if you develop a chronic illness *after* purchasing it. All in all, the tax preference for employer-provided health care serves to offset the adverse selection problem that would otherwise plague this market.

A second market failure associated with health insurance is known as *moral hazard*. This means that because health insurance changes the incentives faced by insured people, they may change their behavior in a way that drives up expenditures on health care. For instance, because most people with insurance face a low or zero out-of-pocket price for additional medical services, they may consume extra medical services for which the true social costs (which are higher than their out-of-pocket costs) exceeds the benefits.

Special treatment of health care might also be justified by equity objectives. The itemized deduction for medical expenditures is consistent with horizontal equity concerns; families that experience large, unavoidable, out-of-pocket medical costs arguably have a lower ability to pay taxes than other families with the same income. Furthermore, vertical equity concerns might motivate a desire to redistribute resources to the poor in the form of health care. In some cases, voters may prefer redistribution for a specific meritorious (in their eyes) purpose such as health care over redistribution of resources that can be used for any purpose, meritorious or not. Moreover, low-income uninsured people often end up getting care anyway—for example, through uncompensated care from a public or charitable hospital. This creates a kind of moral hazard problem by reducing the incentive to buy insurance. While such care does help the poor who receive it, both the benefits and the costs are distributed in a capricious fashion. For instance, whether a particular uninsured individual has access to uncompensated care is largely a matter of luck, and the costs may be borne through higher insurance premiums for certain people or perhaps lower compensation for health-care providers who happen to be willing to work in certain areas. Arguably, a more systematic government policy could provide a fairer and more efficient way of helping the poor than uncompensated care.

Although solid rationales can be found for some form of government intervention in the market for health care, the current tax treatment of



health insurance is poorly designed for addressing these problems and may even make some of them worse. To the extent that helping low-income people afford health insurance is the goal, we now go about achieving that goal all wrong. The tax exclusion of employer-provided health insurance provides no help at all to people whose incomes are too low to be subject to income tax, and it provides very large amounts of help to those high-income people who have high marginal tax rates.

The tax exclusion of employer-provided insurance is also inefficient because it provides an incentive not only to buy health insurance in the first place (which may be desirable) but also to buy expensive health insurance (which is not). Consider an employee who faces a 20 percent marginal tax rate. If an employer wants to give that worker \$80 more in net-of-tax compensation, the firm would pay only \$80 to grant that compensation in the form of a better health insurance policy but \$100 to provide compensation in the form of a higher cash salary. This creates a strong incentive for employers to offer, and employees to prefer, much more generous health benefits than otherwise, which exacerbates the moral hazard problem considerably, leading to some wasteful expenditures on medical care.

For these reasons, eliminating or capping the exclusion from tax of employer-provided medical benefits is often a feature of proposals for health care reform as well as tax reform. Many economists believe that if we want to help people afford health insurance, providing a voucher or credit of a fixed amount would be much more sensible than having a tax exclusion. This would give an equal dollar amount of subsidy to everyone and would eliminate the incentive to buy expensive policies.

In 2007, the Bush administration proposed a plan that would replace the exclusion for employer-provided health insurance, the deduction for self-employed health insurance, and the deduction for out-of-pocket medical expenses. In their place, any taxpayer with a health insurance plan that meets certain requirements would get to exclude exactly \$7,500 (if single) or \$15,000 (if married) from both income taxes and payroll taxes. As long as the taxpayer gets at least some qualifying health insurance, he or she gets the full exclusion. The amount of the exclusion would never change regardless of the actual cost of the insurance premium. Aside from the fixed exclusion, compensation used to pay for health insurance would be subject to both taxes. So the discrete decision to buy at least some qualifying insurance would reduce your



taxes by several thousand dollars, but spending additional amounts on insurance would never reduce your taxes. This would remove the inefficient incentive to buy expensive insurance. Unlike current law, the tax benefit could be used by anyone who gets insurance and would not be limited to just those with employer-provided insurance and the self-employed. This would offer at least some help to many of the currently uninsured. The plan is essentially like a voucher, except that it is worth more to people in higher tax brackets.³⁸

Capping the exclusion or offering a tax credit, voucher, or fixed exclusion for having health insurance would not solve all the problems in the health-care market. For instance, none of these plans addresses the adverse selection problem, and special care would need to be taken in designing such a voucher plan to avoid weakening the incentive for employers to provide group insurance plans, or that problem could become even worse. But a carefully designed voucher-style program for health insurance still has many advantages over the current tax exclusion.

Charitable Contributions

Under current law, contributions to qualifying charitable organizations are deductible from taxable income for those who itemize deductions. Let's subject this provision—which is in effect a penalty on those who are not charitably inclined—to our two-tiered test. First, are charitable contributions an indication that the contributing family is less well off than their income would suggest? Some would argue yes—that people who make charitable contributions are sacrificing some of their own well-being for moral purposes. However, these are voluntary contributions. The contributors must be getting some satisfaction from the act of giving, otherwise they would not have done so. Perhaps they are motivated by the “warm glow” they feel when they help others. So their level of well-being is really no different than that of people with equal incomes who don't give to charity. If that is the case, we cannot invoke an ability-to-pay justification.

Test number two is whether there is something inherent about charity that justifies subsidies to encourage charitable giving. In this case, the answer is arguably yes, as charitable contributions involve positive externalities. For one thing, the beneficiary of the contribution is better off. Second, charitable contributions may provide benefits to people who do not make contributions themselves. For example, many non-contributors might feel better knowing that homeless people are



provided with food and shelter by charities. But in these cases, people have an incentive to free ride on the contributions of others; economists call such a situation a *public goods problem*. Charitable giving still occurs because some people are motivated by the satisfaction they receive from giving, but because of free riding, an inefficiently low amount of giving is done. To the extent this occurs, an extra incentive from the tax code may lead to a higher level of charitable activity that makes everyone better off. Of course, another approach would be for the government to provide the public goods directly, instead of relying on private contributions. However, some things that are arguably public goods would tend to be underprovided by the combination of government and unsubsidized charity. For example, the U.S. government could not contribute directly to religious institutions because of the constitutionally mandated separation of church and state, and some public goods that are valued only by a minority of the population might not be provided by a majority-rule government. Some argue, in addition, that private charitable contributions are a more efficient and less intrusive way of financing public goods such as aid to the poor than is government intervention.

The charitable contribution deduction is not without its costs, however. In 2007, it reduced income tax revenues by about \$50 billion. It adds to complexity and recordkeeping requirements. A broad array of activities qualify as “charity,” so there’s no guarantee that the gifts will go to activities that really deserve to be subsidized by the tax system.³⁹ For example, only a small portion of subsidized giving goes to help the poor; much of it goes to higher-education and cultural institutions, which may or may not deserve subsidy. Contributions can also be difficult to monitor, so unfortunately some inequity arises from abusing the system.⁴⁰ Finally, evidence suggests that the deduction may have only a moderate impact on the amount of giving, so some of the implicit subsidy may serve as a reward to giving that would have occurred anyway.

On balance, some form of incentive for charity may be justified. Even so, it is difficult to justify why the rate of subsidy should be tied to the donor’s tax rate and be zero for taxpayers who do not itemize or who have incomes below the filing threshold—the pattern of subsidy that the deduction for giving creates. More appropriate might be a fixed credit per dollar donated rather than a deduction, so that we are subsidizing all charitable contributions at the same rate and not subsidizing more heavily the contributions of affluent taxpayers.⁴¹



State and Local Government Taxes and Bonds

Two important preferences in our tax code have to do with state and local governments. The itemized deduction for tax payments to state and local governments costs the federal government about \$49 billion in revenues. The exclusion from federal income tax of interest on state and local government bonds costs an estimated \$37 billion.⁴² Are these features justifiable?

One potential rationale for the deductibility of state and local taxes is that families that have the same income but live in high-tax states and municipalities have less ability to pay and therefore should pay fewer federal taxes than families who live in low-tax places. The flaw in this argument is that many of the people who live in high-tax states presumably benefit from a higher level of public expenditures. Why should someone who chooses to live in a low-tax state and make do with fewer government services be penalized for that choice? Of course, the relationship between state and local taxes and the benefits from public services is certainly not one-to-one. So the ability-to-pay argument against deductibility has some merit, although it is limited.

Another problem with the deduction is that it provides an inefficient subsidy to state and local spending. The cost to (itemizing) taxpayers of a dollar of such spending is less than a dollar because part of the cost is shifted to taxpayers in other states. On the margin, this may encourage state and local governments to undertake projects that would not meet taxpayer approval in the absence of the tax incentive and that therefore use up resources that would be more efficiently used for other purposes. Moreover, because it is linked to itemized deductions, this subsidy is larger for governments that have relatively more high-income (and therefore high-tax-rate, itemizing) residents.

Such a subsidy for state and local expenditures might be desirable if those expenditures provide benefits that spread beyond the state's or municipality's borders. In this case, some expenditures that are worthwhile from the country's perspective might not be enacted by a state or local government. This argument certainly does not apply for many expenditures, such as garbage collection or municipal swimming pools. It makes more sense for primary education, on the grounds that it builds an "educated citizenry," which benefits all Americans. Although, this argument has some intuitive appeal, it is hard to prove, and in any event, it applies only to a subset of what state and local governments do and therefore does not justify the general deductibility of state and local taxes. Nor does it justify effectively giving a larger subsidy to



more affluent communities where the residents have higher tax rates.

On balance, if tax reform results in the elimination of deductions for state and local taxes, this would not be a major loss. Although these deductions have some merit as an adjustment for ability to pay and perhaps as an encouragement of certain worthy public expenditures, it also involves significant inefficiency, unfairness, and complexity. The alternative minimum tax disallows the deduction for state and local taxes, so unless the AMT is reformed, the number of people who benefit from this deduction will dwindle over time.

Some of these same arguments apply to the other major income tax provision related to state and local government—the exclusion of interest on state and local bonds from taxable income. In this case, there is no ability-to-pay rationale because the decision to buy bonds issued by municipal governments is entirely voluntary. The main effect of the interest exclusion is to subsidize debt-financed expenditures in the states and municipalities because it enables these governments to borrow money at a lower interest rate than otherwise. High-tax-rate investors are willing to accept the lower interest rates because they pay no tax on the interest. The increased demand by investors drives up the prices of these bonds—or, in other words, pushes the yields on these bonds down toward the after-tax rate of interest offered on similar, but taxable, bonds.

Not all of the benefits of the interest exclusion go to state and local governments, however. If all potential buyers of these bonds had the same tax rate—say, 20 percent—the interest rates on the bonds would end up being about 20 percent lower than on other bonds, so the purchasers would gain little or no benefit. Their *implicit* tax (because they have a lower interest rate) would equal the *explicit* tax on other investments. But in a system of graduated rates and both taxable and nontaxable potential purchasers, to sell all the bonds that state and local governments want to issue, the rate of interest must be high enough to attract not only those taxpayers with the highest tax rate but also many investors with lower tax rates. This means that taxpayers in the top tax brackets benefit because they can invest at a higher after-tax rate of return than otherwise.

Given that the case for subsidizing state and local expenditures is shaky, the interest exclusion seems hard to justify. Because a substantial portion of the subsidy represents a windfall to very high-income people, it is even harder to justify. An alternative to both the exclusion



of bond interest and the deductibility of state and local tax payments that would avoid these problems would be for the federal government to provide direct subsidies to these governments for the intended class of expenditures. This idea has always been resisted by state and local governments, largely due to their fear that once the tax preference becomes a straightforward appropriation, it could more easily go onto the budget chopping block. This is a good example of how the tax system can sustain implicit subsidies that would probably not survive as equivalent stand-alone programs.

The Standard Deduction and Rough Justice

To this point, we have addressed the most important itemized deductions. Only about one-third of all taxpayers deduct these expenses, and these are predominantly affluent families. This is because all taxpayers are offered the option of bypassing the itemizing process and instead receiving a *standard deduction*, which varies only by marital status. In 2007, the standard deduction amounted to \$10,700 for a married couple filing jointly, \$5,350 for a single taxpayer, and \$7,850 for a head of household.

The standard deduction makes sense because it would not be cost-efficient for the IRS to have to monitor and occasionally audit the deductions claimed by the 87 million or so people who file tax returns but do not now itemize, not to mention the cost in time and expense of these taxpayers having to keep track of their expenses. By having a standard deduction, however, the tax system loses its ability to finely differentiate among taxpayers with differing abilities to pay. As it stands now, two otherwise identical families, both with \$30,000 of income and both taking the standard deduction, owe the same tax even though one family has incurred \$5,000 in medical expenses and the other hasn't. Although in principle the tax system allows extraordinary medical expenses to reduce income subject to tax, in practice we settle for "rough justice" by differentiating tax liability only when relatively large sums of money are involved.

If some or all itemized deductions end up being retained, a larger standard deduction could still simplify taxpaying for many people by cutting down on the number of itemizers. This could save substantial administrative and compliance costs. It would mean settling for even more rough justice, but the tradeoff might be worth it. One simplifying change that would cost no revenue would be to couple a reduction or elimination of personal exemption allowances for the adults in a family



with a corresponding increase in the standard deduction. Similarly, if some itemized deductions are retained in a flat tax, the family allowance could be treated as a large standard deduction.

Conclusion

This chapter has addressed the policy issues that arise in contemplating any major tax reform. As with most contentious policy choices, the contemplated changes often require a balancing among the desirable characteristics of a tax system. The next two chapters examine specific proposals for overhauling the tax system. All of these proposals involve some combination of the three elements discussed in this chapter—a single rate (or at least low marginal rates), a consumption base, and a clean base. Each of these three elements of tax reform is conceptually and practically distinct from the others so that a reform could achieve any or all to varying degrees.