Table of Contents

Introduction

During most discussions of tax policy, much of the public's attention is focused on next year's taxes. Often the public debate is centered around a set of tax distribution tables that purportedly show which income group will bear the burden of various taxes. While this debate is interesting, it obscures what is probably a much more important issue of tax policy—how a tax system will affect the nation's long-run economic growth.

Because even small differences in growth compound over time, a nation's standard of living can change dramatically because of small changes in annual economic growth.

Because even small differences in growth compound over time, a nation's standard of living can change dramatically because of small changes in annual economic growth. During the late eighteenth century and for much of the nineteenth, for example, citizens of Great Britain enjoyed the highest living standard in the world. The twentieth century gradually ended that long reign, as year after year the British economy grew slightly more slowly than the United States economy. By 1997, per capita GDP in Britain measured in terms of purchasing power was just 69.8 percent of that in the United States.

Taxes affect economic growth in two ways. First, they transfer resources from the private to the public sector. Here society has to decide on the optimal allocation of resources between the two sectors.

Secondly, they affect the incentives to supply factors of production. A nation's economic growth is a function of its ability to use its capital, labor, and technology to transform resources into goods and services. Taxes affect the level of each of these resources and the manner in which each of them is employed.

A nation's capital stock represents individuals' decisions to save and invest rather than to consume. Taxes that diminish the return on investment dissuade people from saving as much as they would in the absence of the taxes, ultimately lowering the nation's capital stock. Similarly, taxes on labor alter individuals' labor-leisure choices and reduce the amount of labor supplied to the labor market. Likewise, by altering the rewards to innovation, the tax system not only affects the rate of investment in new technologies but affects the level of technology employed in the production of goods and services.

This study provides the reader with an introduction to how the current federal income tax system influences individual behavior with respect to saving and labor supply decisions. The first section examines the effect of income taxes on individuals' decisions to save and invest. Here it is shown that the current federal income tax system generally distorts saving decisions and reduces the amount of U.S.-owned capital with potentially serious consequences for future income growth.

The second section focuses on the effect of taxation on an individual's labor supply decision. Since there are only 24 hours in a day, individuals face a fundamental trade-off between working to make more money and enjoying leisure time. This trade-off is made in terms of income used to purchase goods and services.

A nation's economic growth is a function of its ability to use its capital, labor, and technology to transform resources into goods and services. Taxes affect the employment of all of these.

Taxes affect this trade-off decision by reducing the incentive to work, and they have

their greatest effect on this decision at the point at which an individual is just considering whether to work an additional hour or take an additional hour of leisure. Whether the decision is how much to save or how much to work, taxes are most influential at "the margin." Hence the importance of "marginal tax rates."

Effect of Marginal Tax Rates on Savings and Investment

A nation's economic well-being is greatly influenced by its rate of savings. When individuals forego current consumption in favor of saving, they perform several important functions. First, saving allows individuals to allocate their earnings optimally over time. This is important because in the United States and western countries individuals typically work during only part of their lives. During retirement they maintain their lifestyles by consuming their savings and the income—interest, dividends, and capital gains—earned on their previous saving.

When individuals save out of current income, the savings typically are injected into the financial system, which makes them available to some individual or company needing to make an investment in plant and equipment. The user of these funds is able to acquire them because the lender or equity holder expects to earn a particular after-tax return. When this return is received, some fraction is then paid out to the original saver in the form of dividends, interest, or capital gains. In short, saving allows an individual to lay a claim on future income resulting from investments in productive plant and equipment.[1]

This section presents an overview of how income taxes affect individuals' decisions to save. First the life-cycle model developed by Nobel Laureate Franco Modigliani is presented. Initially this model is used to illustrate how individuals would make their consumption-savings decisions in the absence of income taxes. Such a baseline scenario is useful since economists frequently judge a tax system by the extent to which it distorts choices away from those that would be made in an untaxed market. One that does not distort the market is described as neutral.

A proportional income tax levied on two different tax bases is then introduced into the analysis. In the first case, the tax is levied only on wage income. As we will see, such a tax has no effect on the relative cost of current versus future consumption and is therefore likely to be neutral with respect to saving. Then the tax base is expanded to include capital income. The analysis clearly shows that the inclusion of capital income in the tax base alters the relative cost of current versus future consumption and therefore individuals' saving behavior, and is therefore not neutral with respect to saving. After looking at the effect of a proportional tax system on an individual's propensity to save, we will examine the effects of a progressive tax system levied on wage and capital income. As we will see, such a system may be even less neutral with respect to saving than a flat income tax.

The Life Cycle Model

In the United States and other developed economies, individuals typically work only part of their lives. During retirement they maintain their lifestyles by consuming their savings and the income it generates, whether in the form of dividends, interest, or capital gains. Saving allows people to continue to consume long after they have stopped working, effectively smoothing consumption over their lifetimes.

An individual's saving decision is formally illustrated in Figure 1. Here a person's current consumption, C_0, is given on the vertical axis. Consumption in future periods, C_1, is

[1] An important issue is the relationship between domestic saving and domestic capital formation. In an economy closed to international capital flows, the amount of money saved in a given period determines the amount of investment that businesses can make in plant and equipment during the same period. In such an economy, domestic saving is therefore the source of the capital formation necessary to increase wages and salaries.

Most modern economies are not totally closed. Like the U.S. economy, they are partly or wholly open to international capital flows. Thus, if U.S. saving is insufficient to fund the economy's demand for capital, then the U.S. will be able to make up some or all of the deficiency in domestic saving by importing saving from abroad. In an open economy, there is at most a weak link between domestic saving and domestic investment.

presented on the horizontal axis. For purposes of illustration we will assume that the individual earns wage income during the current period and that he consumes his savings (both principal and interest) during the future period. One could think of the current period as representing the individual's working years while the future period represents his retirement.

One option for the individual would be to consume all of his income during the current period. In this case he would consume at point A in Figure 1. Here he would use all of his income to finance current consumption and he would have nothing left for the future.

Since the individual knows that he must provide for the future, he is unlikely to want to choose point A. Instead, he will save a fraction of his current income to finance future consumption. Fortunately for savers, the market for investible funds provides them with a place where they can invest these funds and earn a positive rate of return, which we will simply refer to as interest. Use of this market opens up all of the points along line AB to the individual. Since interest is earned on invested funds, the slope of line AB is equal to $-1/(1 + r)$.[2] As such, the slope of the budget constraint AB represents the so-called opportunity cost of a dollar consumed, that is, the rate of trade-off between a dollar spent in the current period and the dollar plus the interest earned in the future. If the individual decided to consume at point D, for example, he would consume c_{OD} in the current period and save an amount equal to $(A - c_{0D})$. This sacri-

Figure 1
An Individual's Consumption Choices Over Time

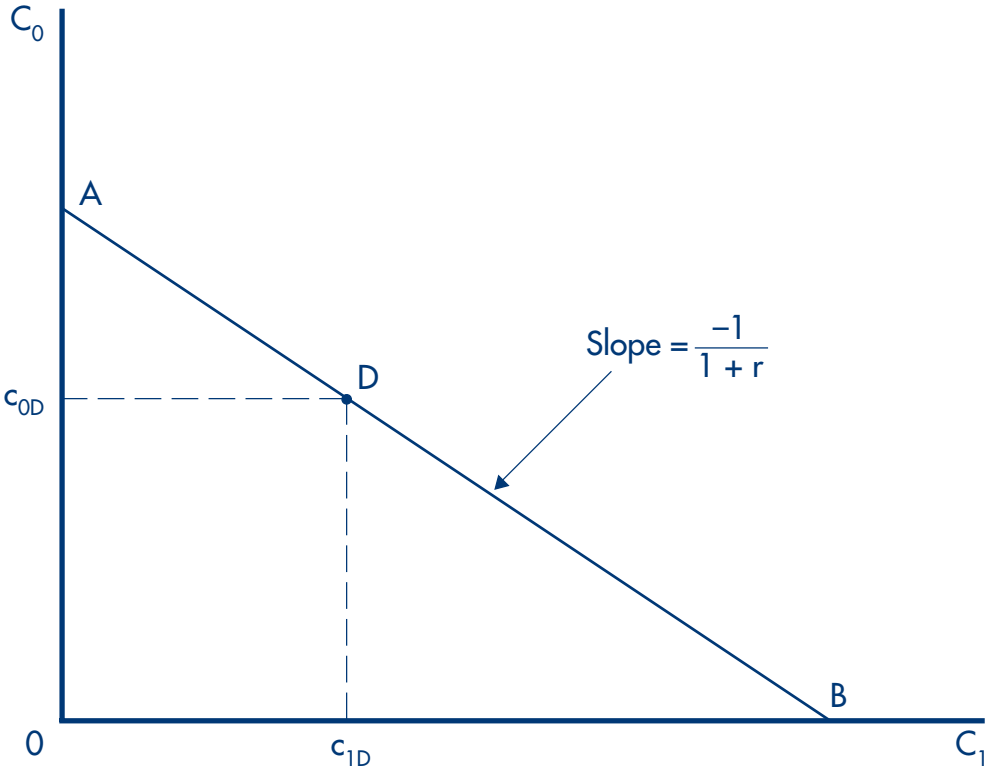

[2] In this model r is the "expected real after-tax rate of return," or the inflation adjusted after-tax rate of return that individuals expect to receive on their savings.

fice and the interest earned on these funds would allow him to consume at c_{1D} in the future period.[3]

Where the individual decides to consume on line AB is determined by his preferences. The concept of "preferences" is both intuitive and crucial. Everyone has preferences. Some prefer small cars, some prefer large cars. Some people prefer white clapboard houses, some people prefer red brick. For any two activities or items, there are combinations for which an individual is indifferent between one possibility and another. For example, consider the trade-off between going to the movies and going out to dinner. Suppose we start with the well-being an individual has over going to the movies three times each month and going to dinner three times each month. If an individual were faced with the prospect of going out to dinner one less time in a month, to remain indifferent between this option and the original option of doing each three times, how many more movies would he or she have to be able to see? The answer to this question yields a new combination of movies and dinners. The individual is indifferent as between the new combination and the original choice of three movies and three dinners.

Figure 2
An Individual's Consumption Choice Between Today and Tomorrow

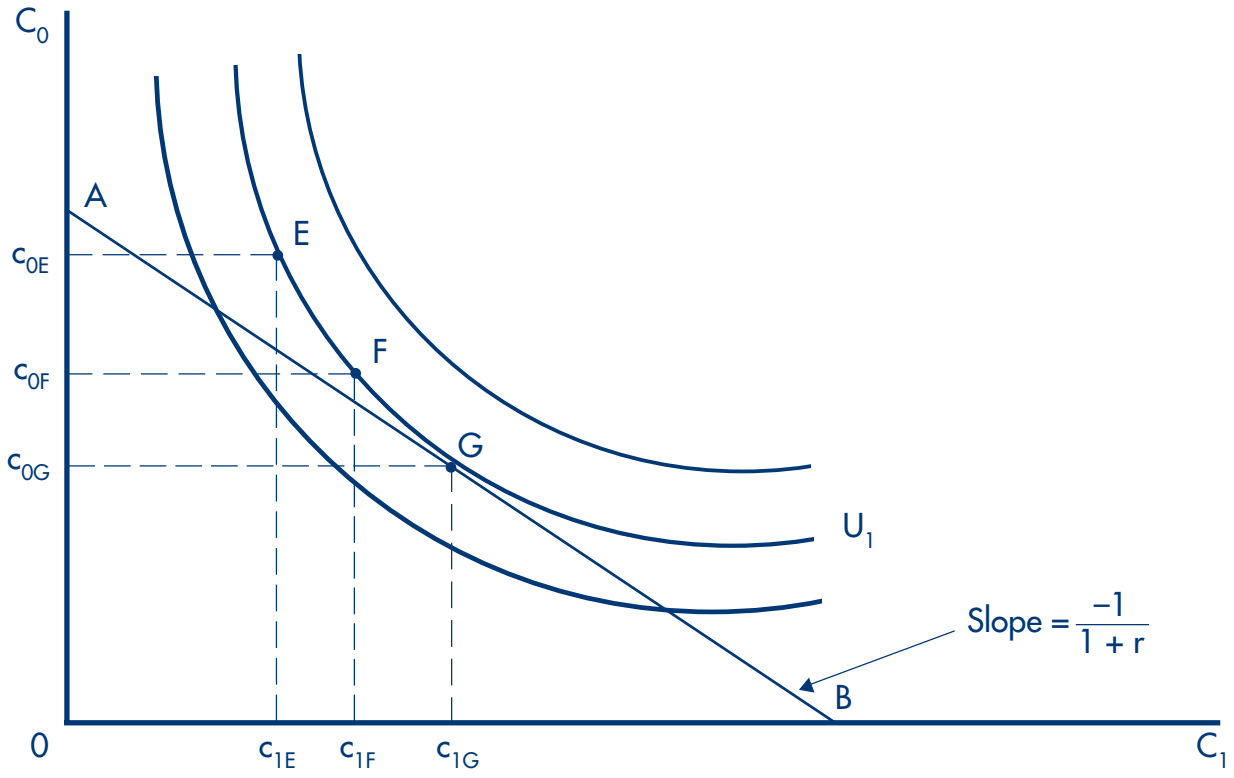

In Figure 2 a set of indifference curves has been added to Figure 1. These curves graphically illustrate the individual's preferences for various levels of well-being.[4] The meaning of

[3] Note how the concept of opportunity cost is illustrated in Figure 1. In the graph the slope of the budget constraint AB is 1.0/1.5. This means that when the individual gave up ($c_{0A} - c_{0D}$) in current consumption, i.e., his savings, he was able to get c_{1D} in future consumption, or 50 percent more. Therefore, every dollar spent in current consumption has an opportunity cost of $1.50 in future consumption.

[4] Economists refer to such well-being as utility.

these curves is quite intuitive. Every point along a curve corresponds to a different combination of current and future consumption that yields the same level of well-being. In Figure 2, for example, point E where the individual consumes c_{0E} during the current period and c_{1E} in the future period yields the same level of well-being as does point F, where the individual consumes c_{0F} in the current period and c_{1F} in the future. Likewise, all of the other points on indifference curve U_1 provide the individual with this same level of well-being.

Figure 3
The Effect of a Proportional Income Tax on Wages on an Individual's Consumption Choices Over Time

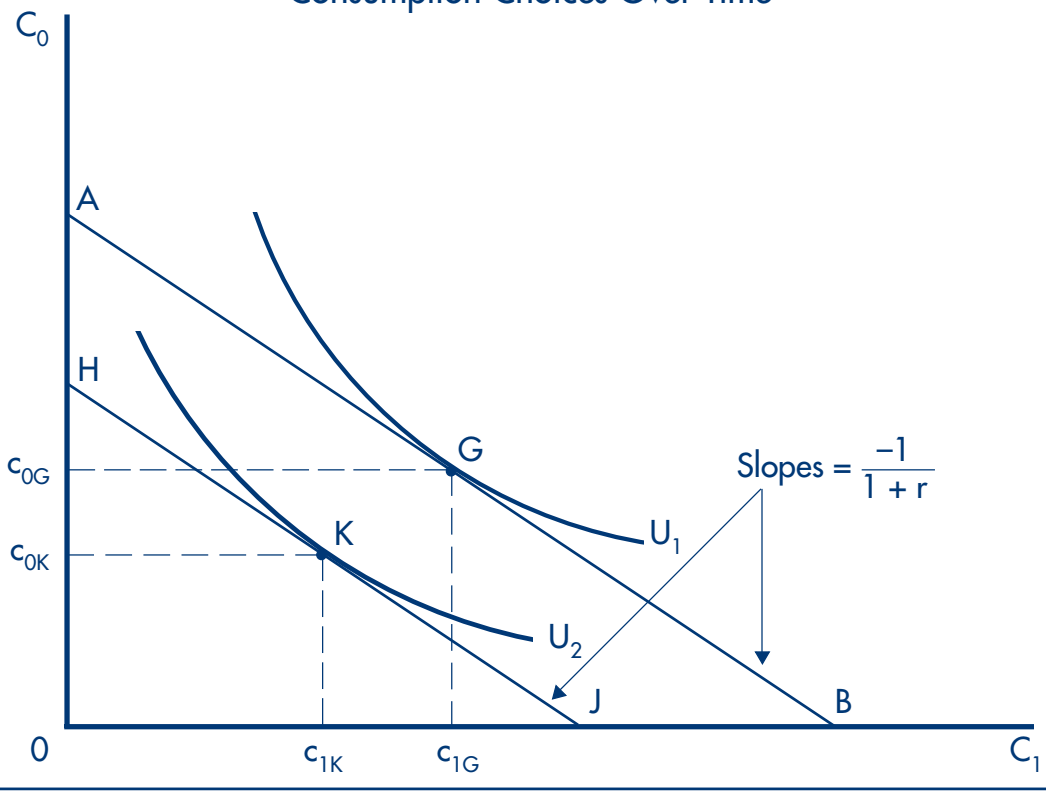

Three related features of indifference curves deserve special attention. First, note that in order to hold a given level of well-being constant, it is necessary that the individual be made to trade off consumption in one period for consumption in another period. Second, note that the curvature, or slope, of the indifference curve represents the rate at which the individual is willing to trade current for future consumption. Finally, since indifference curves further from the origin allow individuals to consume more in one period without necessarily consuming less in another period, they are preferred to those closer to the origin. As we will see, however, lifetime income places limits on an individual's consumption.

Optimizing Consumption

The individual seeks to optimize his lifetime consumption subject to his budget constraint. Such an optimum is achieved at point G in Figure 2 where indifference curve U_1 just touches the budget constraint AB. Note how this point of contact between the indifference curve and the budget constraint occurs where the rate at which an individual *is willing* to trade current for future consumption is the same as the rate at which he *is*

able to trade current for future consumption. At point G the individual consumes an amount equal to c_{0G} during the current period and saves an amount equal to $(A - c_{0G})$. This will allow him to consume an amount equal to c_{1G} in the future.

Introduction of a Proportional Income Tax

Recall that the individual whose budget constraint and preferences were illustrated in Figure 2 earned two types of income during his lifetime. During the current period he earned wage income, some of which he consumed and some he saved. During future periods the individual consumed what he had saved of this wage income plus the interest, or capital income, he had earned. This is an important distinction since taxing each evokes a different behavioral response.

Taxing Wages

A tax on wage income will shift the budget constraint in a parallel fashion toward the origin. However, since it does not affect the relative costs of present and future consumption, such a tax will not change the slope of the budget constraint. Therefore, any change in saving will be dictated by the loss of income brought about by the imposition of the tax and not by a change in the relative costs of current versus future consumption.

A wage tax is neutral in the sense that it leaves unaffected the relative cost of current and future consumption. At the very least, a wage tax is likely to reduce the absolute amount of saving. It may also reduce the proportion of income that is saved because an individual's propensity to save may vary with income. Intuitively, an individual with a modest amount of after-tax income is likely to save less, all other matters held equal, than an individual with much more income. At a higher rate of tax, an individual's after-tax income will be less, and so he or she may choose to consume a greater share of his after-tax income in the current period.

Taxing Interest

In contrast to a wage tax, a tax on interest income will affect the slope of the budget constraint.[5] To see why, recall that the slope of the budget constraint represents the opportunity cost of current consumption in terms of future consumption. When an income tax is levied on interest income it means that money saved for future consumption will have less after-tax buying power. As a result, the opportunity cost of consuming during the current period falls relative to saving for the future. In isolation, we would therefore expect the individual to substitute current for future consumption. However, to some extent this substitution effect is offset by an income effect. An income tax on interest income will lower lifetime income. As a result, the individual must save more out of current income to achieve a given level of future consumption. Therefore, the substitution effect causes the individual to save less after the imposition of the income tax on interest, while the income effect causes him to save more. Which of these effects dominates his decision making is determined by his preferences. Whatever the outcome, however, a tax on interest income almost certainly changes an individual's savings behavior and is therefore not neutral with respect to saving.

Modeling the Wage and Interest Taxes

A case where only wage income is taxed is illustrated in Figure 3. Here an income tax of rate t_1 is levied on the wage income of the individual whose preferences were illustrated in Figure 2. Recall that the individual initially established his optimal trade-off of current and future consumption subject to his budget constraint at point G, at which point he was consuming c_{0G} during the current period and saving $(A - c_{0G})$. This would have allowed him to consume c_{1G} in the future. Imposition of an income tax on wage income lowers the individual's income during the current period

[5] Note however that it will not change the C_0 intercept since if the individual consumed all of his income during the current period he will be unaffected by the tax on interest income.

from 0A to 0H. This causes the budget constraint to shift inward toward the origin from AB to HJ. After imposition of the tax the individual optimizes his consumption on the indifference curve U_2 at point K, at which point he consumes c_{0K} during the current period and saves (H – c_{0K}). This allows him to consume c_{1K} in the future.

The case where capital income is also taxed is illustrated in Figure 4. Here the tax is levied on both the wage and capital income of the individual. As before, the person was initially optimizing his balance of current and future consumption subject to his budget constraint at point G. Imposition of an income tax has two effects. As in Figure 3, the tax on wage income shifts the individual's budget constraint from AB to HJ. In isolation this would result in the person optimizing at point K. However, when interest income is also included in the tax base there is another effect. Taxing interest changes the slope of the budget constraint from $-1/(1 + r)$ to $-1/(1 + r(1 - t_1))$.[6] This is illustrated in Figure 4 by the pivoting of the budget constraint from HJ to HL. Given this new constraint and the individual's preferences optimization will occur at point M, at which point the person would consume c_{0M} in the current period and save (H – c_{0M}). This would allow them to consume c_{1M} in the future.

The following text is relatively difficult, and so the casual reader may want to jump to

Figure 4
The Effect of a Proportional Income Tax on Wages and Capital Income on an Individual's Consumption Choices Over Time

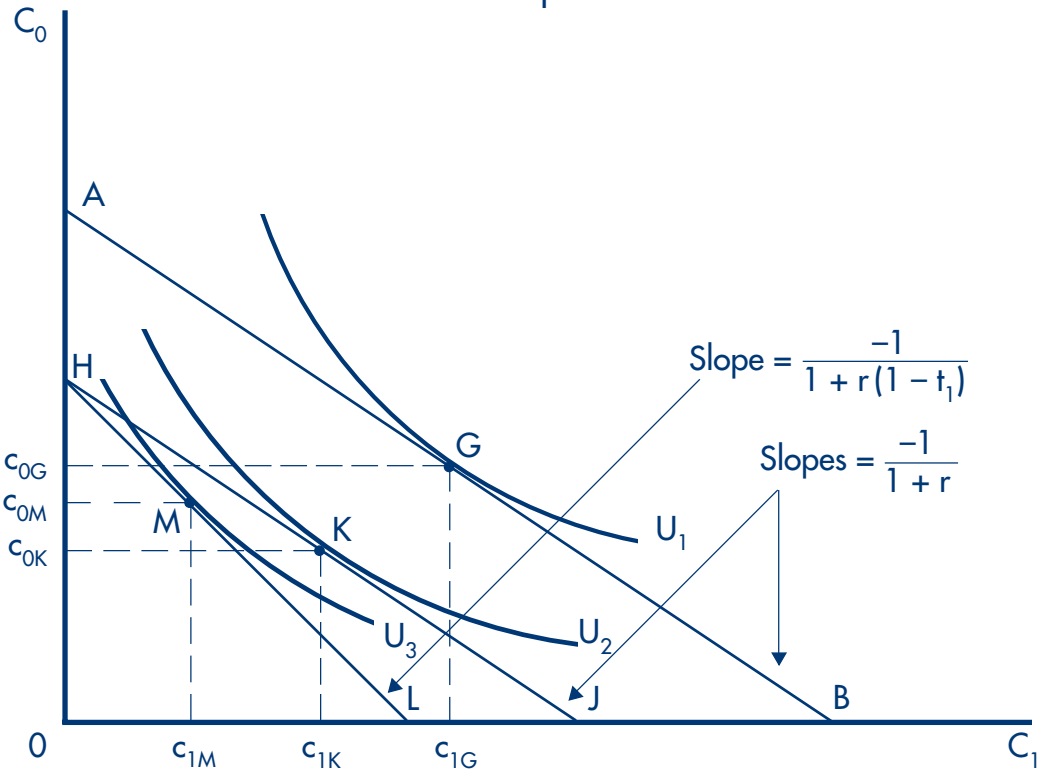

[6] The presentation assumes the rate of interest does not itself change in response to the imposition of a tax on interest income. In reality, the rate of interest increases significantly, evidencing a "tax wedge" between the pre-tax and after-tax rates of interest. This tax wedge reflects an implicit marginal tax rate paid by the marginal saver. The balance of the presentation reflects the common instance in which the marginal tax rate of the potential saver exceeds that implicit in the tax wedge.

the end of this section.The material is important, however, and is included for completeness. Figures 3 and 4 are presented so that the reader can clearly see how the imposition of income taxes on different tax bases influences behavior. Comparison of the two graphs is somewhat inappropriate, however, since the tax base is larger in the latter case.

Figure 5
Comparison of a Wage Income Tax and a Revenue Neutral Wage and Capital Income Tax

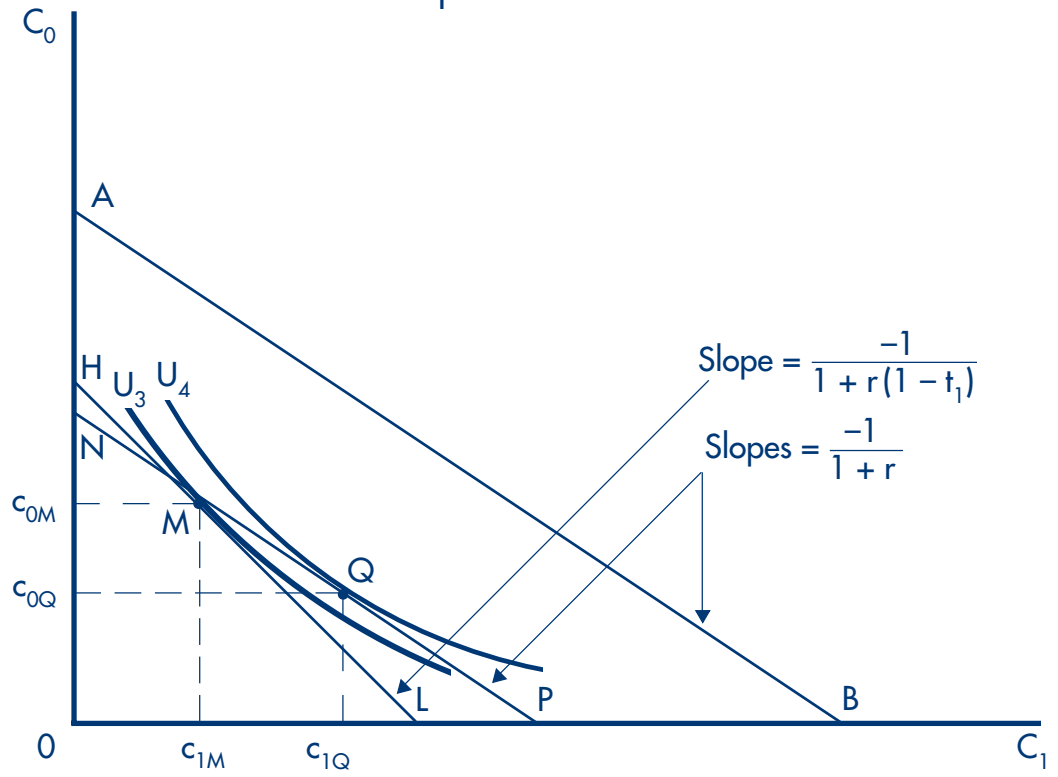

Figure 5 presents the case where the tax systems each raise the same amount of revenue. Here interpretation of point M is the same as in Figure 4. Point Q, on the other hand, represents the individual's preferred choice if the income tax were only levied on wage income. Since the interest accrued during the future period is not taxed in this case, a slightly higher tax rate, t_2, is required to make it revenue neutral. This reduces the individual's current income by AN which is slightly more (HN) than in the case where both wage and interest income were taxed (AH). However, since there is no tax on interest in this case, the slope of the budget constraint NP is equal to that of the original constraint AB. In this case the individual optimizes his balance of current and future consumption at point Q. At this point he would consume c_{0Q} in the current period and save ($N\text{-}c_{0Q}$). This would allow him to consume c_{1Q} in the future. In this case it appears as though the individual would save slightly more under a wage income tax than under an equal yield wage and capital income tax. However, as stated above, this is by no means a rule. It depends on the preferences of each individual.

Introduction of a Progressive Income Tax

So far we have assumed that all taxpayers are subject to the same proportional income

tax. However, the current federal income tax system is progressive, meaning that people who earn more income pay a higher percentage of their income in income taxes. In large part this progressivity is achieved through the use of a rising marginal tax rate structure. Currently, the federal individual income tax system has five statutory rates. As one's income grows, it becomes subject to successively higher rates. In 1998, for example, individuals paid 15 percent on their first $25,350 in taxable income. On the next $36,050 they paid 28 percent. Taxable income between $61,400 and $128,100 was subject to a rate of 31 percent. A 36 percent rate was paid on the next $150,350 of income and all income above $278,450 was taxed at the 39.6 percent rate. Taxable income is calculated by subtracting exemptions and deductions from adjusted gross income.

The progressive, statutory marginal tax rate schedule and the many factors that go into calculating taxable income such as the standard deduction and the Earned Income Tax Credit cause individuals' average and marginal tax rates to differ. In 1998, for example, an individual who took the standard deduction and had an adjusted gross income of $32,301 would have been in the 28 percent tax bracket. However, his average tax rate would have been only 11.8 percent. This is because the standard deduction and personal exemption would have made the first $6,950 of his income free of federal individual income tax. The next $25,350 of his income would have been subject to the 15 percent rate. Therefore, only $1 of his income would have actually been subject to the 28 percent rate. Overall the marginal tax rate structure creates a progressive income tax system where the effective tax rate rises with income.

Aside from the variety of statutory income tax rates that a taxpayer faces depending on his taxable income, that tax rate varies according to whether the taxpayer is married or single, how many children are claimed as dependents, whether there are long-term capital gains to declare, etc. The tax rate that applies to the taxpayer's last dollar of taxable income is the marginal tax rate.

The marginal tax rate is generally quite different from an average or an effective tax rate. A taxpayer's average income tax rate is the ratio of the amount of taxes paid to taxable income. A close synonym for "effective" is "actual," but the exact meaning depends on the context. An effective marginal tax rate may differ from a marginal tax rate because the taxpayer may be in an income range in which he is subject to a phase-out of some exclusion or deduction. An effective average tax rate may differ from an average tax rate because some measure of income other than taxable income is used. For example, the Joint Tax Committee typically calculates the effective average tax rate as the ratio of taxes paid to a constructed measure of "economic income."

Each of these measures of the tax rate has its appropriate time and place. In assessing the economic effects of a tax system, the appropriate tax rates are the marginal or the effective marginal tax rates.Taxpayers face their highest tax rate, which is their marginal tax rate, at exactly the point at which they are most likely to be influenced by taxes. Thus, the marginal tax rate and not the average or the effective average tax rate is of most importance in determining the effect of taxes on individual behavior.

Figure 6 illustrates the effect that a rising effective tax rate on wage and capital income has on an individual's propensity to save. In the figure the individual whose budget and preferences were illustrated in Figure 4 is subject to progressively higher effective tax rates t_1, t_3, and t_4. For purposes of illustration the individual's income is held constant. The figure shows that an increase in the tax rate has two effects. First, it shifts the budget constraint toward the origin. Secondly, it changes its slope, making it steeper as the tax rate rises. This second effect shows that the substitution effect becomes more pronounced under a progressive income tax system. In this case the individual becomes increasingly less inclined to save as his effective tax rate rises. However, as stated before, because the substitution and income effects work in opposite directions it is not possible to assert whether overall savings will decline or rise.

Responses to Higher Effective Tax Rates

Except for very low-income individuals, most taxpayers have a high degree of flexibility in choosing how much to save and how much to consume in a period. In the short run, of course, one's ability to reduce consumption may be limited because people often enter into contractual arrangements such as car loans and mortgages that require a certain amount of spending. Over the longer term, however, one's flexibility can be significant and tends to grow with one's income and wealth.

Whatever the circumstances, as the foregoing discussion implies, the saving versus consumption decision is essentially a decision made "at the margin." Similarly, a taxpayer's savings-versus-consumption decision is going to be affected primarily by his marginal tax rate. Specifically, whatever capital income is earned next year on this year's saving, that income will be subject to the taxpayer's marginal tax rate. The higher the marginal tax rate, the lower will be the expected after-tax return to saving, and the more pronounced will be each of the income and substitution effects.

Real World Complications

The theoretical model of individual behavior presented here is fairly simple in its essentials. Individuals can consume all or only part of their current income. Their choice as to how much to save will be determined at the margin by their preferences and how they view a dollar of income consumed today versus a dollar plus some return consumed tomorrow. The relevant tax parameter in this discussion is the marginal tax rate, since that is the rate that will be applied tomorrow to the income earned on today's saving.

Figure 6
The Effect of a Progressive Income Tax Levied on Wage and Capital Income

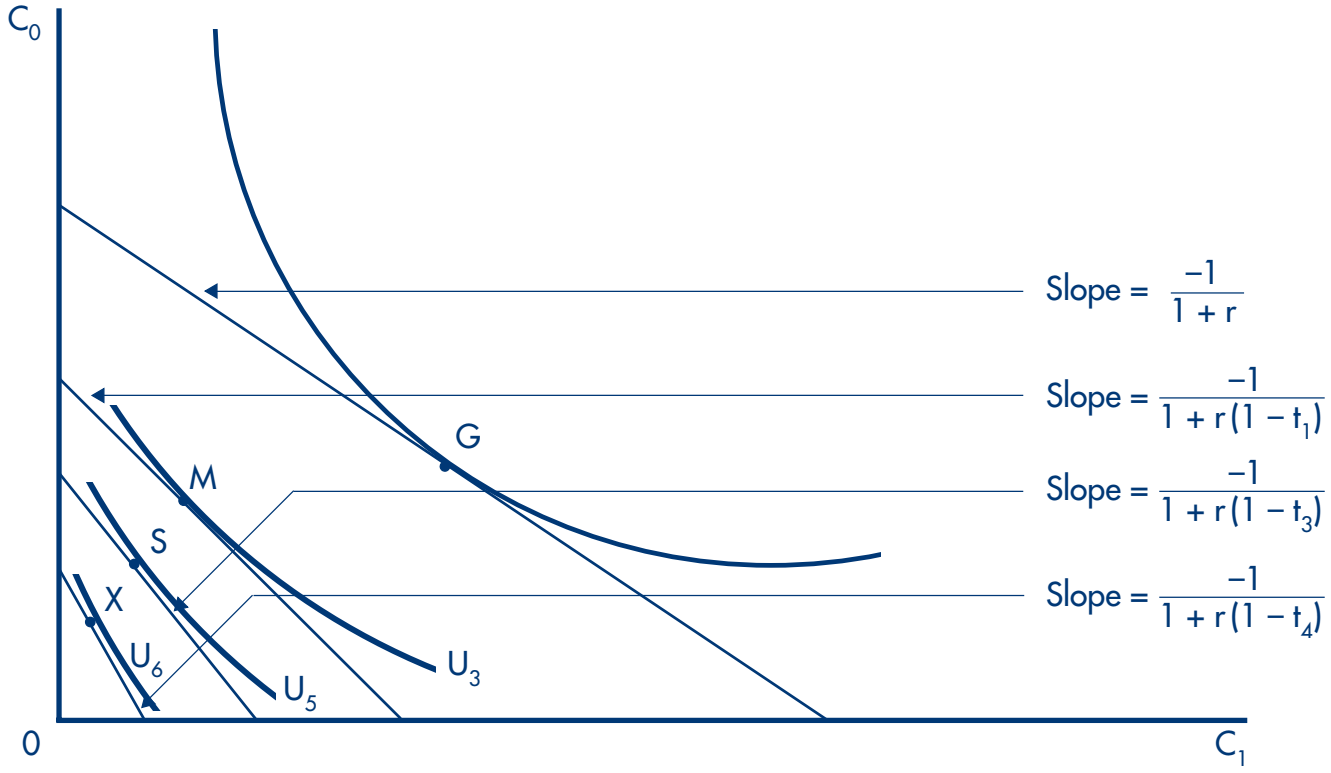

Wealth and Age

While this simple model is correct as far as it goes, the real world poses a number of complications. For example, taxpayer responses to different marginal tax rates vary significantly depending on one's income or wealth. Everything else held equal, as one moves up the income scale, the share of income that is saved tends to increase. While there are obvious exceptions, people tend to value the next dollar's worth of consumption slightly less than the previous dollar's consumption. This being the case, it may be that a taxpayer's responsiveness to a higher marginal tax rate declines as his income rises.

The possibility of bequests may also alter the effects of marginal tax rates on saving behavior. A taxpayer facing a high marginal tax rate may choose to make a bequest, either to a family member, friend, or charitable organization, rather than continue to pay a high rate of tax on the income accruing to savings. In this case, the amount transferred is neither saving nor consumption.

Also, how a taxpayer reacts in his saving behavior to a particular marginal tax rate may vary according to his stage of life. For example, a taxpayer who is currently saving nothing out of current income may be unlikely to save even less through borrowing if the tax rate is higher. This might apply to low-income taxpayers or to individuals in retirement. Similarly, taxpayers who are just starting a family are likely to have less disposable income and therefore may be relatively immune to marginal tax incentives or disincentives to save. Other taxpayers not in one of these special situations may react strongly to a higher or lower marginal tax rate because they have the flexibility to do so.

The Complexity of the Tax Code

The tax code itself greatly complicates the question of how marginal tax rates affect saving behavior. For example, the tax treatment of long-term capital gains differs from that of income. Except in times of high inflation, long-term capital gains generally receive a lighter tax treatment than do other forms of capital income. Historically, capital gains have faced a lower statutory tax rate. Taxpayers also have a limited capacity to offset their capital gains with capital losses. Finally, taxpayers are allowed effectively to erase their capital gains liability at death, though they may face a steep estate tax instead.

Aside from capital gains, the tax code displays a wide array of special saving vehicles, all of which make the income tax operate more like a simple wage tax. Individuals can choose "front-loaded" or "back-loaded" Individual Retirement Accounts. Many taxpayers still have old-style company pensions, while others have employer-based 401(k) or 403(b) retirement plans. In addition, the tax code now has Medical Savings Accounts to help individuals self insure their health expenses and the code has Education Savings Accounts to help parents save for their children's college expenses.

With so many tax-preferred saving vehicles, for many taxpayers it is unclear exactly what their marginal tax rate is on a marginal dollar of saving. Take for example the case of an individual with a 401(k) retirement plan. These plans allow workers to save up to a certain amount of current income by contributing the income to the plan in pre-tax dollars. Also, all earnings on the amounts saved are tax-deferred. However, the plans have limits on the amount of current income that can be contributed in a given year and they have strict limits and penalties on early withdrawals. If a taxpayer does not contribute the maximum allowed to his or her 401(k), one could conclude that the 401(k) has completely eliminated any income tax disincentive to this individual's saving, in which case the taxpayer's marginal tax rate is whatever he or she expects it to be at the time the money is withdrawn. In this case, with a 401(k) plan, the income tax operates like a wage tax. This may not be the case, however, if the taxpayer feels constrained by the limits on early withdrawals.

Because of the limitations and penalties on early withdrawals, tax-preferred retirement plans generally are poor substitutes for short-term saving or saving for emergencies. Thus, even if one has a 401(k) plan and does not

make the maximum contribution allowed, the relevant marginal tax rate is the regular income tax rate if one's marginal dollar of saving would not be for retirement but for future emergencies.

On the other hand, if the taxpayer makes the maximum allowed contribution in a given year, then the relevant marginal tax rate is whatever that rate would be if there were no special retirement plan since that is the relevant tax rate on a marginal dollar of saving.

Empirical Investigation

With so many variations in taxpayer situations and so many complications added through the tax system to an individual's marginal saving-versus consumption-decision, it is almost surprising that some of the empirical literature is able to discern any relationship at all between saving behavior and marginal tax rates.

The inability of economic theory to state precisely how taxes affect savings has led to a great deal of empirical investigation. Until the late 1970s much of this research tended to show that saving rates were largely unresponsive to changes in the after-tax rate of return.[7] This could mean either that tax policy has little effect on saving or that the income and substitution effects brought about by policy changes roughly offset each other. In 1978 Michael Boskin released a study in which he argued that such findings were incorrect and that, when properly measured, the real after-tax rate of return had significant effects on saving. These findings led to a debate that continues to this day.[8]

The problem plaguing research in this area is that many of the important variables are unobservable and therefore must be estimated. Recall that a central feature of the life-cycle model is its assumption that saving behavior is governed by the expected real after-tax rate of return. There are several problems associated with calculating such a figure. First, since there are many types of assets in the economy there is no single rate of return that can be studied. This has forced researchers to construct composite, economy-wide rates of return. Similarly, since individuals' saving decisions are presumably made after considering real rates of return, estimates must be made of individuals' inflation expectations. Since no one knows for sure how such expectations are formed conjectures must be made.[9] Different assumptions about how composite rates of return are constructed and about how inflation expectations are formed have led to vastly different conclusions about how the real after-tax rate of return affects saving. Some argue that estimation problems make research in this area futile. Von Furstenberg (1981), for example, argues that the real net after-tax rate of return effect on saving is virtually impossible to ascertain since it is almost impossible to estimate such a rate.

An even more complicated issue arises in assessing the after-tax returns facing savers. Every investment option bears some degree of risk: risk about the ability of the asset to generate the expected income, risk about the forecast of inflationary expectations, risk about the future tax treatment to which the income may be subjected, risk that the institution where the funds are placed for investment may fail, etc. Measuring these risks as they face individual savers, or as they are perceived by individual savers is virtually impossible. Even more problematic, there is no reason to presume these risks are perceived by savers to be constant over time. When prices are stable, for example, it is easy to be confident in one's forecast of future inflation. But when inflation is high or variable, the uncertainty surrounding one's forecast will rise dramatically.

Another issue is that different savers face different investment options. Wealthier indi-

[7] See Wright (1969), Weber (1976), and David & Scadding (1974). A review of prior research is also contained in Boskin (1978)

[8] For a good summary of this literature see Bernheim (1996). Some researchers have even called into question the validity of the life-cycle model itself. See Kotlikoff and Summers (1981). For a contrary view see Hurd (1986).

[9] See Boskin (1978) and Blinder and Deaton (1985).

viduals often have access to investment options such as initial public offerings unavailable to small savers. Also, as noted above, different savers may be subject to different tax treatments on their marginal dollars of saving. For example, one saver may work for a business with a generous tax-deferred pension plan whereas the next saver is self-employed. Or consider two savers who work for the same firm. One saver takes full advantage of the company's tax-deferred pension and has additional, fully taxed saving. The other saver has a minimal level of saving in the pension and no other saving. These two savers may be identical in every other respect, but the after-tax return on their marginal dollar of saving will be very different.

Policy Implications

The current federal income tax system taxes both wage and capital income. As we have seen, such a system is unlikely to be neutral with respect to the choice between current and future consumption. While empirical evidence is inconclusive as to whether tax policy is artificially depressing the rate of saving, the fact that Americans appear to save less than citizens of other industrialized countries has led many to believe that Americans are jeopardizing their future by their failure to save. Many feel that a low saving rate is leading to an insufficient provision of resources to fund the retirement of a graying population.

Policymakers have responded to such concerns in two ways. The first is by enacting numerous saving incentives such as Individual Retirement Accounts, Keogh, and 401(k) plans. Others argue that such measures do not go far enough and feel that the current system should be replaced by either a wage tax or some type of consumption tax. Since consumption would always be taxed at the same rate, regardless of when it occurs, a properly constructed consumption tax would be more nearly neutral. All of the major tax reform proposals currently under consideration, such as the Armey-Shelby Flat Tax and the national retail sales tax, would effectively tax current consumption and leave saving untaxed. The aim of all of these plans is to introduce a great degree of neutrality into the tax system.

Effects of Marginal Tax Rates on Labor Supply

Introduction

Policymakers need to be aware of the consequences of tax policy on the labor supply to implement efficient tax policy. Good policies help boost national output that in turn boosts household income. The relationship between tax policy and household income is especially important when considering the labor supply. After all, a primary determinant of income is the number of hours worked. Inefficient tax policy creates disincentives to working and lowers the standard of living for many American families.

This section examines the extent to which marginal tax rates influence the level of labor supplied in the economy. The economic theory of the labor supply, while providing a useful framework, is not conclusive in its predictions. Ultimately, the empirical evidence on the labor supply must decide, first, if changes in marginal tax rates also cause movement in the labor supply and, second, if the movement is significant. This section concludes with a brief look at some recent research in this area.

Theoretical Issues

People work to earn money to purchase goods and services. On the other hand, people also enjoy leisure time. Since there are only 24 hours in a day, people face a fundamental trade-off between work and leisure.

In order to make this trade-off, individuals must consider the opportunity costs of a given choice. The opportunity cost is a measure of the foregone options after a choice is made. In this case, if individuals decide to pursue more leisure, they would forego income by working less. Therefore, when discussing the trade-off between work and leisure, the gained or lost income measures the opportunity cost.

The trade-off between labor and leisure is constant as long as a person's wage does not vary with the number of hours worked. If the wage is $10 an hour, for example, then working one hour more earns $10 more, whether the increase is from 4 hours to 5, or from 8 hours to 9. While the opportunity cost of working more or less is therefore linear, the nature of the trade-off as perceived by a worker is not linear. Someone working 16 hours a day, for example, is likely to value one hour of leisure much more than someone working eight hours a day. Similarly, someone working two hours a day is likely to place less value on an hour of foregone leisure than is someone working eight hours a day. In short, the more leisure one has, the less one values an extra hour of leisure, while the less leisure time one has, the more one values an extra hour of leisure.

Figure 7 graphically depicts how a single individual resolves the work-leisure trade-off and thereby determines his labor supply at a given wage. The vertical axis represents the gross income earned while the horizontal axis represents the number of hours in a day that can be used for work or leisure.[10] The straight line connecting the two axes represents the budget line that plots the opportunity costs of leisure in terms of income. As drawn, the budget line reflects the assumption that the hourly wage does not vary with the number of hours worked.

The U-shaped line represents the work-leisure trade-off that yields a constant level of satisfaction to a given worker.[11] For example, in Figure 7, points A, B, and C are all equivalent in terms of the satisfaction level. Point A yields $50,000 in income, but leaves only 3 hours a day for leisure. At the other extreme,

[10] Note, the leisure axis only accounts for 16 hours of the day, the remaining 8 hours are considered to be used for sleeping hence they are unavailable for purposeful work or leisure.

[11] Economists formally refer to this as the indifference curve. The indifference curve here represents all the proportional trade-offs between leisure and income that yield the exact same utility.

point C only yields $20,000 in income, but leaves 12 hours a day for leisure. However, points A and C are not attainable because they do not fall on the budget line. As a result, point B is the "optimal" choice because it affords the highest levels of income and leisure—given the budget constraint, shown by line B_1, and the preferred work-leisure trade-off, shown by line U_1. At point B, the marginal cost of an extra hour of leisure, or the opportunity cost of an extra hour of leisure, is exactly equal to the marginal benefit of an extra hour of leisure as perceived by this particular individual.

Assessing how an individual changes his work-leisure trade-off in response to a change in his wage is complicated because there are two identifiable reactions—the income and substitution effects—that work in opposite directions. Ultimately, the two effects can lead to either a decrease or an increase the number of hours worked in response to a change in the wage rate.

Figure 7
Concepts Used for Labor Supply Analysis

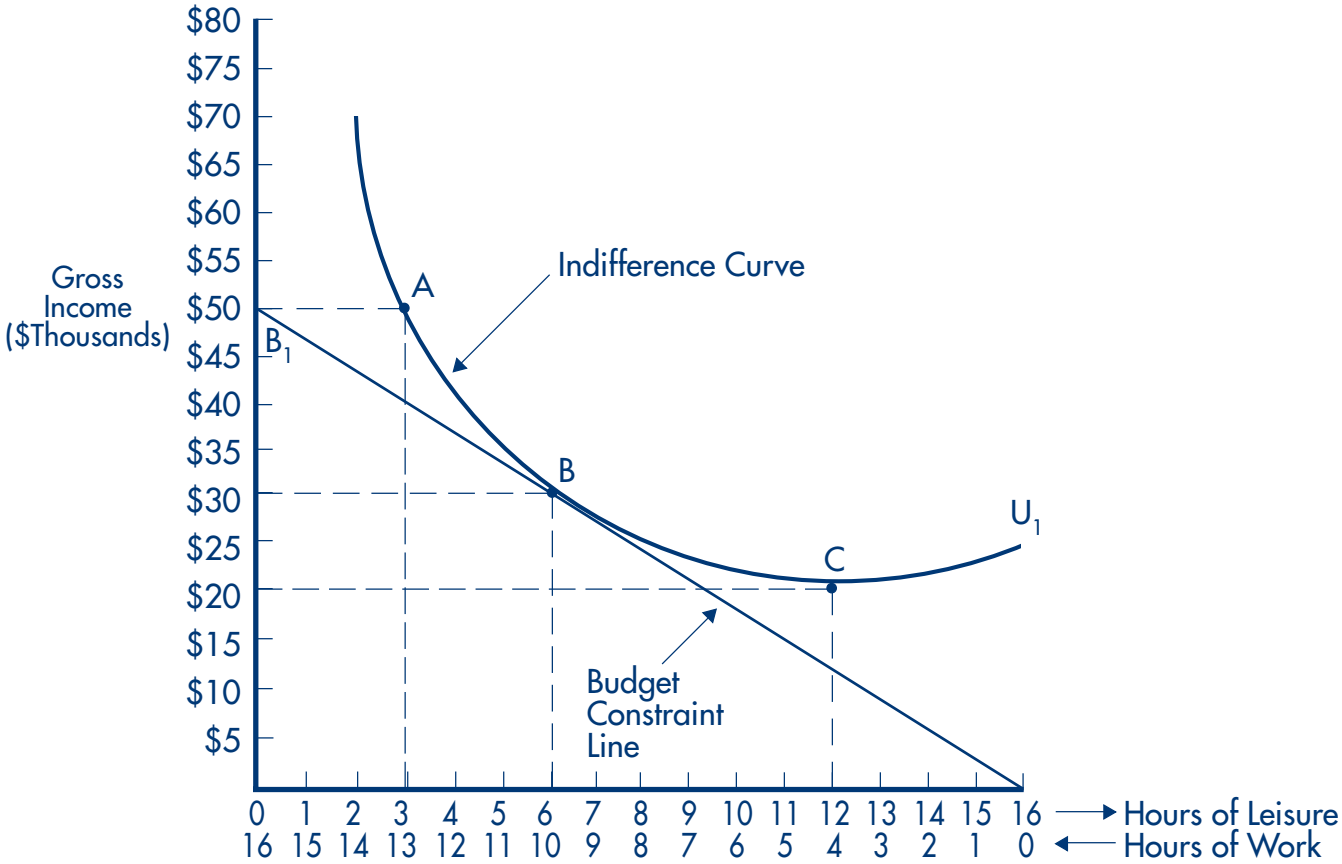

Changing Work Hours Following a Wage Increase

As shown in Figure 8, a wage increase boosts potential income and is represented graphically by shifting the budget line up from B_1 to B_2.[12] The individual is now better off because he earns more income per hour

[12] Note, the budget line pivots on the horizontal axis because a change in income does not affect the number of total hours available for work or leisure.

worked. The individual now faces two choices. First, he can decide to work less and enjoy more leisure. This is known as the income effect. Second, he can work more to reap more fully the income benefits of the higher wage. This is known as the substitution effect.

In this case, the individual decides that the higher earning opportunity is too good to pass up. As a result, he may only increase leisure by 1 hour from 6 to 7 because of the income effect (A to B). However, swamping the income effect, the individual decides to increase work by 3 hours because of the substitution effect (B to C). Overall, the final decision represented by point C shows a 2-hour increase in work from 10 to 12 hours and a corresponding 2-hour decrease in leisure from 6 to 4 hours.

Alternatively, the individual could decide that he can afford an extra 3 hours a day in leisure. In Figure 9, the movement from point A to point B illustrates the income effect. However, the higher wage is tempting so he decides to reduce his increase in leisure by 1 hour. The movement from point B to point C illustrates the substitution effect. The final decision is therefore represented by point C which shows a 2-hour decrease in work from 10 to 8 hours and a corresponding 2-hour increase in leisure from 6 to 8 hours.

Figure 8
Increase in Income with Substitution Effect Dominating (Work Hours Increase)

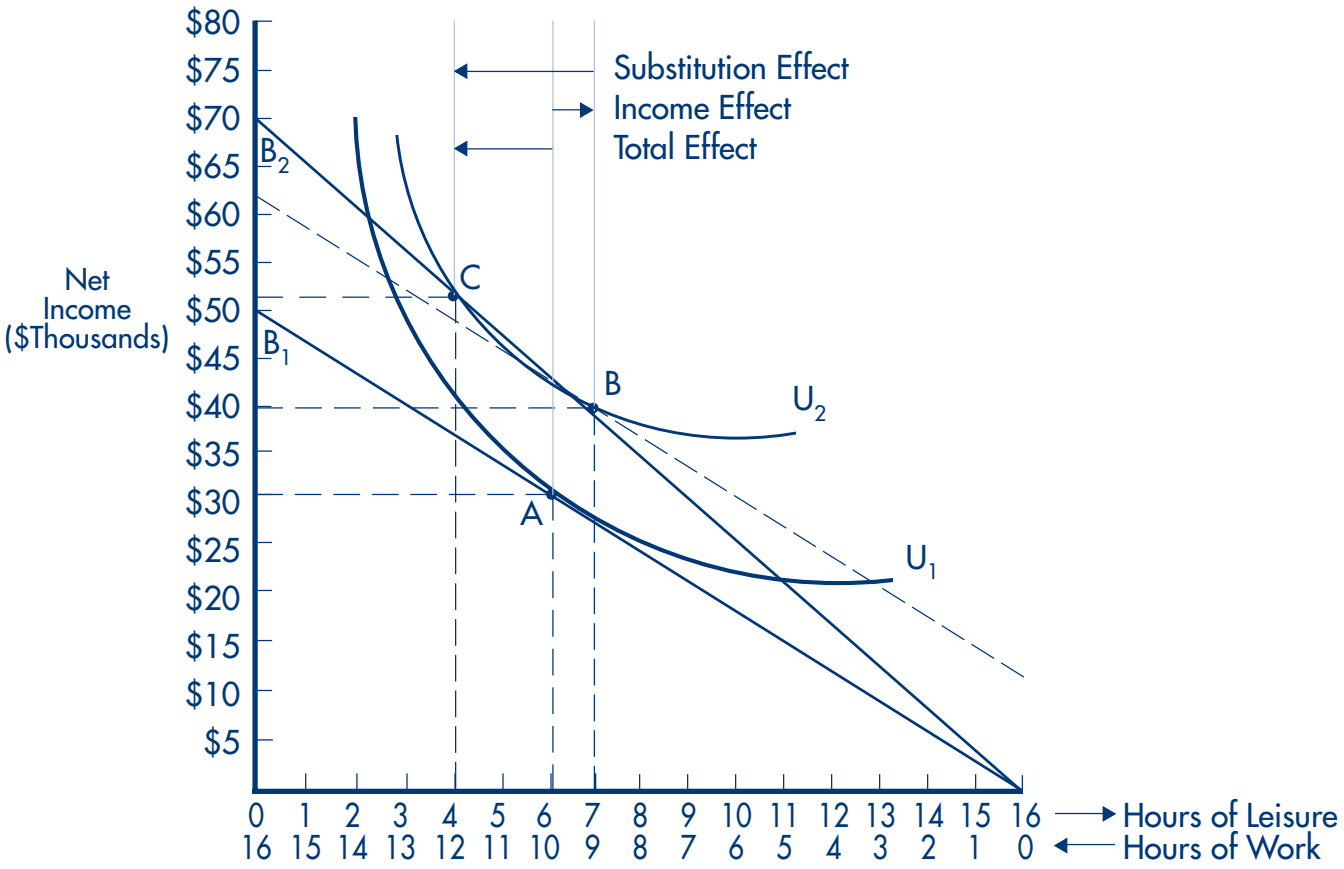

The Labor Supply

An issue of great interest is how the supply of labor changes over time in response to

changing wages. The matter must be framed properly, however, to get sensible results. Clearly, the response to a dollar an hour increase from an initial wage of $5 an hour would likely be different than if the initial wage were $15 an hour. A much more useful means of expressing these relative changes is to compare the percentage change in the wage and the percentage change in the labor supply. If one takes the ratio of these percentages—the percentage change in labor supplied over the percentage change in the wage rate—then one has a single statistic capturing all the relevant information. Economists make great use of this statistic and call it the elasticity of labor supply.

A common criticism of this presentation of the labor supply and wage rate is that workers rarely have much freedom to choose the number of hours they wish to work at their going wage. Many jobs, for example, offer no overtime, so the worker cannot increase the hours worked at his or her primary job. It may be possible to take a second job, but then the wage is likely to be lower, in which case the incentive to take the job is reduced, or there may be restrictions on the minimum number of hours that must be worked at the second job. It is also true, however, that at many jobs overtime hours are paid a premium, such as time and a half, so the incentive to work extra hours is even greater.

These concerns about an individual's ability to pick their work hours, while valid,

Figure 9
Increase in Income with Income Effect Dominating (Work Hours Decrease)

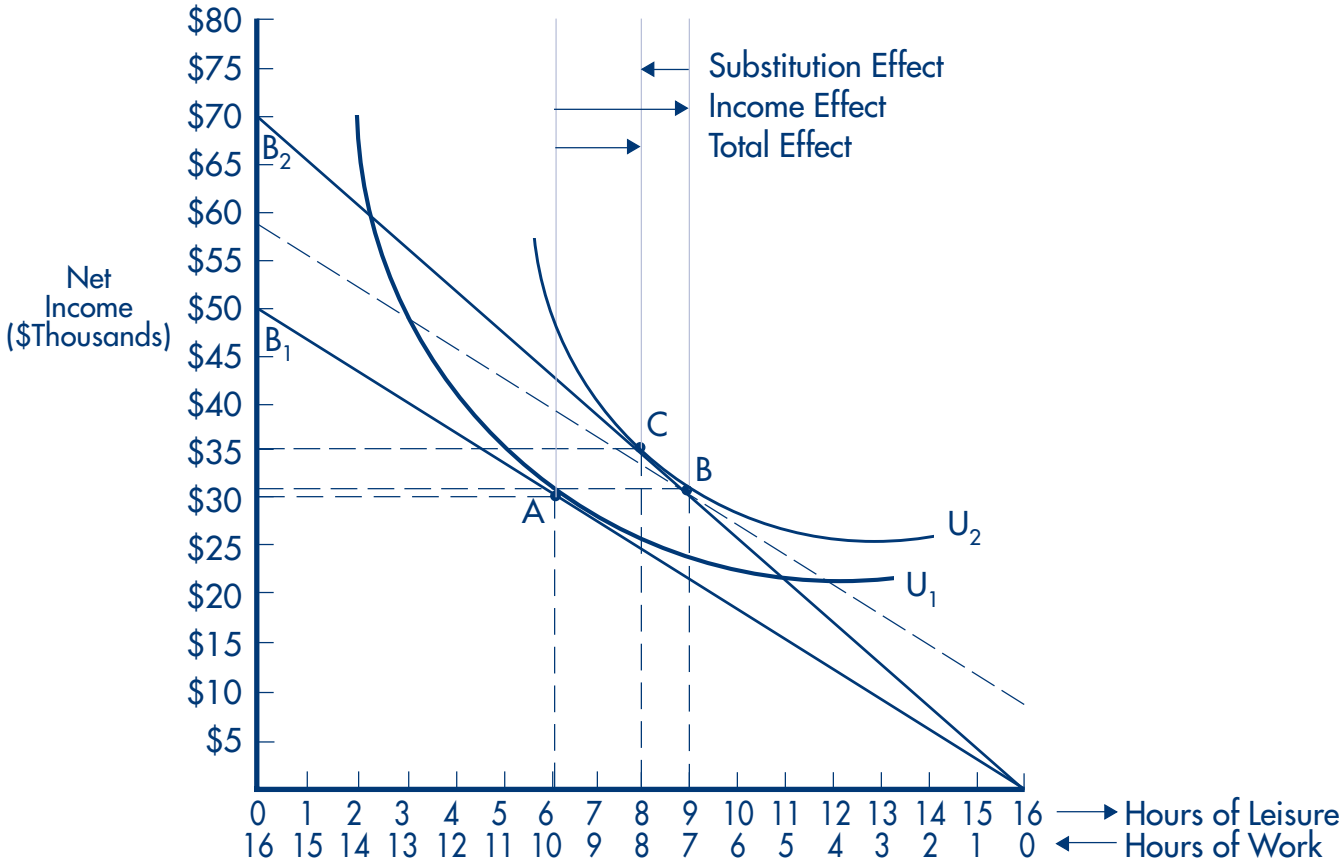

should not be overstated. Given enough time and effort, most workers have the ability to change their work hours to some extent. Particularly in normal economic times, the United States is blessed with a relatively open labor market. Workers can generally change jobs, change locations, and negotiate new terms of employment if given time and a reason to do so.

Effects of Marginal Tax Rates on Income and Labor Supply

The effects of taxation can be incorporated into the work-leisure trade-off through its effects on income. For example, suppose a flat wage tax is imposed. The first point to notice is that unlike Figure 7, Figures 8 and 9 have the vertical axis labeled as net income which is gross income minus taxes ($I_n = I_g - T$), where the amount of tax paid is some constant fraction of total income. The linear budget line is representative of a proportional tax, i.e., the tax has a single marginal tax rate. Therefore, the increase in after-tax income shown in Figures 8 and 9 could be attributed to a reduction in the tax rate.

Notice that the effect of reducing the tax is to increase the true opportunity cost of leisure in terms of work in that, after taxes, the amount of income foregone for an additional unit of leisure is greater than is the case when the tax is higher. Facing this new marginal return to work, the individual adjusts his or her desired amount of work and leisure until the marginal value of another unit of work is exactly equal to the after-tax marginal income received from working.

In the examples described thus far it is assumed that a single rate or proportional wage tax or income tax is imposed. Consequently, the after-tax wage is the same however many hours an individual chooses to work. In this case the marginal rate of tax is always equal to the average rate of tax, defined as the total amount of taxes paid over total wages.

The Effect of a Progressive Rate Structure

In a progressively designed tax system, marginal tax rates increase at higher income levels. For example, the federal individual income tax has statutory tax brackets of 0, 15, 28, 31, 36, and 39.6 percent, with each higher bracket taking effect at a higher level of income than the preceding bracket. In addition, the tax code contains many phase-outs of deductions, exemptions, and tax credits that can create a much higher effective marginal income tax rate than is implied by the statutory rates. For example, the per child tax credit phases out at the rate of $50 for each additional $1,000 in income beginning at $110,000 in adjusted gross income for joint filers.This raises the effective marginal tax rate a half percentage point over the statutory rate.

Table 1
Effect on Income of Hypothetical Progressive Income Tax

Gross Income	Marginal Tax Rate	Additional Tax Paid	Cumulative Tax Paid	Additional After-Tax Net Income	Cumulative After-Tax Net Income	Effective Average Tax Rate
$0 - $10,000	5%	$500	$500	$9,500	$9,500	5.0%
$10,001 - $20,000	10%	$1,000	$1,500	$9,000	$18,500	7.5%
$20,001 - $30,000	15%	$1,500	$3,000	$8,500	$27,000	10.0%
$30,001 - $40,000	20%	$2,000	$5,000	$8,000	$35,000	12.5%
$40,001 - $50,000	25%	$2,500	$7,500	$7,500	$42,500	15.0%
$50,001 - $60,000	30%	$3,000	$10,500	$7,000	$49,500	17.5%
$60,001 - $70,000	35%	$3,500	$14,000	$6,500	$56,000	20.0%
$70,001 - $80,000	40%	$4,000	$18,000	$6,000	$62,000	22.5%
$80,001 - $90,000	45%	$4,500	$22,500	$5,500	$67,500	25.0%
$90,001 - $100,000	50%	$5,000	$27,500	$5,000	$72,500	27.5%

The obvious effect of a progressive tax system is that the amount of taxes paid increases faster than income. Table 1 presents a hypothetical progressive tax system and its effects on net income. In this example, taxpayers face a progressive income tax with 10 equally spaced brackets ranging from 5 to 50 percent. All income is taxable. As shown in the "net income" column, the increasing marginal tax brackets reduce after-tax net income as gross income increases. As a result, the effective average tax rate—which is the percentage of taxes paid to gross income—increases with the rising income.

Steeply progressive tax rates can have a surprisingly strong effect on work effort. As an individual's income rises, he or she will have greater flexibility to choose leisure over labor. Individuals with relatively low amounts of income must use a high percentage of their purchases for necessities. As an individual's income rises, the share of income dedicated to necessities drops and the share devoted to discretionary items increases. For example, at a given amount of income, individuals might be satisfied driving a small "econobox" car. As their incomes rise, they are able to move into more reliable, more spacious, or more luxurious cars.

If an individual with a significant amount of discretionary income faces a high marginal tax rate, he or she must compare the value of the marginal income foregone from working less with the marginal value of additional leisure. In other words, he is effectively comparing the value of a slightly less spacious car with the value of a few more hours a week to devote to his favorite leisure activity.

The implications for the economy in general of steeply progressive tax rates is also an issue. In a market economy, people generally earn an amount of income roughly equal to what they contribute to the economy. Whether one is a star quarterback, a cook, or a computer programmer, one gets out of the system in terms of income what one puts into the system in terms of value. It immediately follows, then, that progressive taxation means that, generally speaking, the most productive participants in the economy face the greatest tax disincentives to using their talents to increase their income and the nation's.

Figure 10 reflects the common view on the theory of labor supply that the substitution effect generally outweighs the income effect. In other words, for reasonable increases in wages most individuals would always choose to work more when wages rise instead of taking additional leisure time. Under this scenario, a progressive tax system is unambiguously harmful because it constrains the labor supply to a lower point than would exist otherwise. A single rate, or proportional, tax system is optimal under these assumptions.

While most workers would respond to a decrease in their after-tax wage with a desire to reduce the number of hours worked, this response is by no means universal. There are individuals up and down the income scale who could have the opposite reaction. For example, a low-income worker may need a certain amount of income to support his family. If the tax rate goes up and his after-tax wage drops, he will probably choose to work more hours to generate the necessary amount of income. Likewise, there are upper-income individuals who have the economic freedom to cut back their hours worked, who would respond to an increase in their tax rate by working more hours. These cases, however, are the exception rather than the rule.

Of course, reality is more complex than theory, and in the case of tax policy one complication is particularly important. Most people receive their labor income as a package of wages and salary plus benefits such as health insurance, dental insurance, and so forth. Unlike wages and salary, which are subject to income tax, most employee benefits are exempt from tax. Recognizing the disparate tax treatment between cash and benefits, workers and employers have gradually shifted an ever-greater share of employee compensation into the form of tax-exempt benefits. Employers are happy to oblige the preferences of their workers because they are allowed to deduct the expenses no matter what form they take.

Obviously, as an employee's marginal tax rate increases, the greater the incentive to shift

out of taxable forms of remuneration into non-taxable forms. The most highly paid members of the workforce therefore have the greatest incentive to shift their pay packages into tax-exempt forms. Clearly, however, despite the high cost of health insurance, there is a limit to how much health insurance an individual needs. The same holds for dental insurance or disability insurance. Further, there are strict limits as to how much a person can put into tax-preferred retirement plans. Given these limitations and given relatively high marginal tax rates, the market has popularized alternative forms of tax-exempt or low-tax compensation forms, such as employee stock options.

An important consequence of this tax-induced shifting of compensation packages is that high marginal tax rates tend to raise less revenue than simple revenue projections would indicate. It also means that reducing marginal tax rates reduces revenue less than expected because workers have less of a disincentive to receive tax-exempt compensation.

Ultimately, the final response of the labor supply to changes in income is theoretically indeterminate. This is due to the fact that the income and substitution effects work in opposite directions. As a result, assumptions must be made about which effect dominates the other. Based on such assumptions, one can only hypothesize whether increasing marginal tax rates reduce or increase the labor supply. The next section will provide a look at recent empirical studies to help determine the most likely result of rising marginal tax rates.

Figure 10
Labor Supply Curve

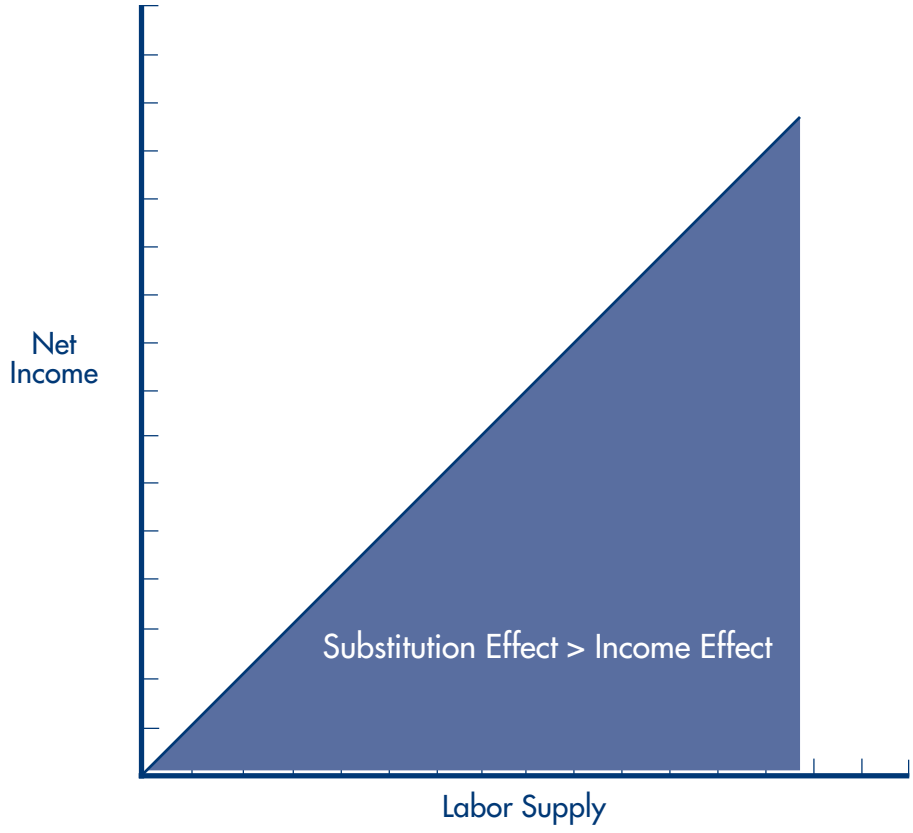

Empirical Evidence

To date, the best empirical studies on the effects of tax policy on the labor supply were written after the passage of the Tax Reform Act of 1986 (TRA '86). Martin Feldstein identifies three reasons for this. First, the magnitude of the tax changes were very broad. Second, many of the changes in the tax code were largely unexpected before 1986, therefore making it possible to compare pre-reform years with post-reform years. And third, government data in the 1980s was qualitatively superior to data generated in the past. For these reasons, this empirical review will limit itself to empirical studies written after 1986.[13]

Despite the improved quality of the data and the recent changes in marginal tax rates that afford a good case study, the real-world effect on the labor supply from marginal tax rates is less obvious than the theoretical model suggests. In fact, changes in the labor supply can arise that are not reflected in simple measures of man-hours worked. For instance, Harvey Rosen has detailed three other major dimensions in addition to the conventional hours worked in a day: lifetime hours of work and timing of retirement; intensity of work effort; and quality of work effort.[14] Therefore in the short-term, changes in the marginal tax rate on work behavior may not be immediately noticeable because they occur at a much later date (such as taking early retirement) or they affect productivity (such as taking more breaks).

While quantifying many of the factors suggested by Rosen remains beyond the reach of economists, reliable data sets do exist that are useful for empirical work on the labor supply. These data sets contain detailed data on taxes and work hours. This data comes from a variety of sources such as the Treasury public-use sample of individual tax returns and the Current Population Survey produced by the University of Michigan.[15]

Although the tax reforms in the 1980s have provided economists with a full decade of empirical data, the difficulties in empirical testing have narrowed these labor supply studies to specific groups. For the most part, they are primarily focused on groups of taxpayers who are most likely to have been the most affected by the decrease in marginal tax rates-such as high-income households and secondary earners (married women).

The Response of High-Income Households

Recent studies have looked at how marginal tax rates affect high-income households which are believed to be particularly sensitive to marginal tax rates for several reasons. First, they are often subjected to higher marginal tax rates. Second, the tax reforms in the 1980s provided them with the largest drop in marginal tax rates. Finally, they also have more opportunities for altering their behavior. Although the recent literature on this topic shows mixed results, the overall evidence points to a slight increase in hours worked as a result of lower marginal tax rates for high-income households.[16]

The Response of Married Taxpayers

Recent studies have also looked at how marginal tax rates affect the work effort of married men and women. Since married men are usually the primary workers in most households, the literature contends that their work hours will respond the least to changes in marginal tax rates. For instance, the primary wage earners are likely to be involved in developing a long-term career out of their work. As such, reductions in work hours or work effort because of higher marginal tax rates would prove to be detrimental to the longer-term goal of professional development.

Married women, on the other hand, would be more responsive for two reasons. First, they often have the choice between working in the taxable formal sector, such as paid employment, or in the non-taxable informal sector, such as rearing children. Second, because their income is generally secondary, their income is

[13] See Feldstein (1995).
[14] See Rosen (1980).
[15] However, even this data has its critics. See Heckman (1993).
[16] See Bosworth and Burtless (1992), Moffitt and Willhelm (1998), and Knieser and Ziliak (1998).

often taxed at a higher marginal tax rate.[17] Third, secondary earners often work part-time. Part-time jobs provide more flexibility not only in terms of scheduling, but also in terms of hours worked, whereas full-time primary employment almost always demands a minimum, 40-hour work week. As a result, secondary earners can react more quickly to changes in the marginal income tax rate. Offsetting these factors is the growing share of women in the work force with professional careers. Professional women are generally subject to the same constraints as professional men in terms of needing to continue working full time to preserve their professional development. Despite the increasing role of professional women, most of the studies done on this issue have confirmed that married women's work hours are significantly more sensitive to changes in the marginal tax rate than are married men's.[18]

Conclusion

In the introduction of this section, two critical questions were asked about the effect of marginal tax rates on the labor supply. In response to the first question, the empirical literature generally finds that higher marginal tax rates do reduce the labor supply and there appear few, if any, circumstances in which a higher marginal tax rate can be expected to increase the labor supply. Thus the evidence supports the proposition that the substitution effect generally outweighs the income effect. As a result, taxpayers respond to cuts in the marginal tax rate by working more hours and to increases in the marginal tax rate by working fewer hours. However, in response to the second part of the question, the magnitude of the response is still being debated.

It appears that a taxpayer's responsiveness to a change in marginal tax rates is very much dependent on his or her particular situation. For example, for single people and for the primary earner in a dual-earner household, the labor supply response to a change in marginal tax rates is very muted, at least in the short-run. This may be due to an inflexibility in labor market arrangements such as employers demanding 40-hour work weeks, for example. It may be due to the fact that these workers tend to be relatively accepting of the offered after-tax wage and thus have little inclination to substitute leisure for labor in the face of a higher after-tax wage. It may be that the income effect and the substitution effect roughly offset one another. Or it may be that all three explanations are important for different subgroups of this part of the labor pool.

Higher income households are also likely to respond somewhat more intensely to changes in the marginal tax rate. For example, the income effect of a reduction in marginal tax rates may be greatly subdued because these individuals may already have a good deal of wealth. Thus, they are more likely to remain in the workforce or to work more hours in the face of a lower tax rate. On the other hand, because they are wealthier they have a greater ability to withdraw from the workforce if the marginal tax rate increases, thereby reducing the marginal cost of additional leisure.

The response of a second earner in a dual-earner family to a change in the marginal tax rate may fall somewhere between the average response of the affluent and that of the primary earner in such a family. A dual-earner family is likely to have a higher level of income than a single-earner family, all else held equal. Thus, they may have greater ability to reduce the total number of hours worked in the face of a higher marginal tax rate. However, many dual-earner families have low rates of current saving, implying they are spending their income as fast as they earn it. Thus, it might be difficult to reduce hours in the face of a marginal tax rate increase, at least in the short run. In the longer run, however, their capacity to reduce their spending may be so significant that the second earner may choose to remain at home.

[17] Secondary income is defined as any income earned in addition to primary income. Oftentimes the addition of secondary income bumps taxable household income into higher tax brackets. As a result, the net increase in household income is smaller for secondary income earners thus making leisure options more attractive to them.

[18] See Eissa (1995), Hausman and Poterba (1987), and Saez (1999).

References

Savings and Investment

Bernheim, B. D. 1996. "Rethinking Saving Incentives." In Alan Auerbach, ed., *Fiscal Policy: Lessons From Economic Research* Cambridge, MA: MIT Press. 1997.

Blinder, A.S. and A. Deaton. 1980. "The Time Series Consumption Function Revisited." *Brookings Papers on Economic Activity II* 1985. pp. 465–511.

Boskin, M.J. 1978. "Taxation, Saving, and the Rate of Interest." *Journal of Political Economy* 86, pp. S3–S27.

David, P.A. & Scadding, J.L. 1974. "Private Savings: Ultra Rationality, Aggregation and Denison's Law." *Journal of Political Economy* 82, pp. 225–49.

Hurd, M.D. 1986. *Savings and Bequests. Working Paper No. 1826* Cambridge, MA: National Bureau of Economic Research. January 1986.

Kotlikoff, L.J. and L.M. Summers. 1981. "The Role of Intergenerational Transfers in Aggregate Capital Accumulation." *Journal of Political Economy* 89, pp. 706–732.

Von Furstenberg, G.M. 1981. "Saving" in H.J. Aaron and J.A. Pechman (eds). *How Taxes Affect Economic Behavior* Washington, DC: Brookings Institution.

Weber, W.E. 1976. Comments by Warren E. Weber in H.J. Aaron and J.A. Pechman (eds). *How Taxes Affect Economic Behavior* Washington, DC: Brookings Institution. 1981.

Wright, C. 1969. "Saving and the Rate of Interest." In A.C. Harberger and M.J. Bailey (Eds.), *The Taxation of Income From Capital* Washington, D.C.: Brookings Institution.

Labor Supply

Bosworth, Barry and Gary Burtless, "Effects of Tax Reform on Labor Supply, Investment, and Saving," *Journal of Economic Perspectives* Winter 1992, Vol. 6, No. 1, pp. 3–25.

Eissa, Nada. 1995. *Taxation and Labor Supply of Married Women: The Tax Reform Act of 1986 as a Natural Experiment. Working Paper No. 5023* Cambridge, MA: National Bureau of Economic Research.

Feldstein, Martin. "Behavioral Responses to Tax Rates: Evidence from the Tax Reform Act of 1986." *American Economic Review* June 1995, Vol. 85, No. 2, pp. 170–174.

Hausman, Jerry A. and James M. Poterba. "Household Behavior and the Tax Reform Act of 1986," *Journal of Economic Perspectives* Summer 1987, Vol. 1, No. 1, pp. 101–119.

Heckman, James J. "What Has Been Learned About Labor Supply in the Past Twenty Years?" *American Economic Review* May 1993 (Papers and Proceedings), Vol. 83, No. 2, pp. 116–121.

Kniesner, Thomas J. and James P. Ziliak. 1998. *The Effects of Recent Tax Reforms on Labor Supply* Washington, D.C.: American Enterprise Institute Press,

Moffitt, Robert A. and Mark Wilhelm. 1998. *Taxation and the Labor Supply Decisions of the Affluent. Working Paper No. 6621* Cambridge, MA: National Bureau of Economic Research.

Rosen, Harvey S. "What is Labor Supply and Do Taxes Affect It?" *American Economic Review* Vol. 70, No. 2, May 1980, pp. 171–176.

Saez, Emmanuel. 1999. *The Effects of Marginal Tax Rates on Income: A Panel Study of 'Bracket Creep'. Working Paper No. 7367* Cambridge, MA: National Bureau of Economic Research.