


Until now, our discussions of capital income taxation have treated saving as a single activity and assumed that any taxes imposed on capital income are applied as the income accrues. In reality, many types of assets serve as vehicles for saving, and the taxes imposed on saving often vary by asset (e.g., debt vs. equity vs. housing) or purpose (e.g., short-term saving vs. retirement saving). Also, the method of imposing taxes on capital accumulation varies, including taxation on accrual of income but also taxation on sale of assets (in the case of capital gains) or death (in the case of estate and inheritance taxes). Heterogeneity of tax treatment means that there are potential behavioral responses to taxation other than simply in the amount saved, and that questions of incidence and efficiency are also more complicated.


Taxation and Portfolio Choice

Under a progressive income tax, some individuals may have higher marginal tax rates than others on capital income. This difference in marginal tax rates may affect not only the level of individual saving, but also the *composition* of saving, in terms of assets. There are two potential reasons why the composition of assets, i.e., the individual portfolio choice decision, may be influenced by taxation. The first is that not all assets are taxed at an individual's normal tax rate. The second, discussed below, is that taxes affect not only expected returns, but also the riskiness of returns. For a variety of reasons, some assets may be tax favored, that is, face a lower rate of tax than the individual's regular marginal tax rate. An important example is assets that deliver their income in the form of capital gains, which are taxed less heavily than other income in most countries. In the United States since 2003, lower tax rates also apply to dividends, so that all income from investment in equity, both dividends and capital gains, is tax favored relative to income from fully taxed assets, such as debt issued by companies or by the federal government.

To understand the impact of the existence of tax-favored assets, suppose there are two assets that are perfect substitutes from an investor's perspective except in the way they are taxed. Income from debt (i.e., interest) is taxed at the taxpayer's full rate of tax, θ , while income from equity is taxed at rate $\lambda\theta$, where $0 < \lambda < 1$. If r^d and r^e are the before-tax returns to debt and equity, then



the relative returns for investors are as shown in this graph. Investors with regular tax rates above θ^* will prefer equity, those with lower tax rates will prefer debt, and those at θ^* will be indifferent. Thus, we'd expect tax-favored assets to end up in the portfolios of high-bracket investors. If we added a third asset that is even more tax-favored, for example municipal debt, which is tax exempt, the picture would look like:




There will now be three groups of investors: those in the highest brackets, above θ^{**} , holding municipal debt; those between θ^* and θ^{**} holding equity; and those below θ^* holding taxable debt.

Note that we have drawn the graphs assuming that the more tax-favored the asset, the lower its before-tax rate of return. This must be true if positive amounts of all three assets are to be held. Otherwise, all investors would prefer tax-favored assets.

This translation of tax-favored status into lower before-tax rates of return is another illustration of tax capitalization, in this case the capitalization of tax benefits. For the marginal investor on the borderline between two assets, the tax benefits of the more lightly taxed asset are just offset by the asset's higher price (i.e., lower before-tax rate of return). How much before-tax rates of return differ depends on the relative asset supplies. For example, if the supply of municipal debt were reduced, then a higher value of θ^{**} , with fewer investors holding municipal debt, would clear the market for municipal debt. Put another way, in this model the tax benefits that are capitalized are those of the marginal investor, whose identity depends on asset supplies. Note that in this model the incidence of taxation comes in two components, through capitalization and direct taxation. For example, individuals who hold equity bear some tax through a lower before-tax rate of return than on debt, and some through their own (favorable) taxation of equity returns. Individuals who hold municipal debt bear taxes only indirectly, through capitalization. This distinction between capitalized taxes and taxes directly paid is important to keep in mind when looking at statistics on tax burdens that reflect only the latter.

So far, our characterization of tax-induced portfolio choice is unrealistic in predicting that each individual will specialize in a particular asset, which is clearly at odds with actual portfolio patterns. This is because assets typically differ in another dimension as well – their risk profiles. Differences in risk tolerance and a desire for portfolio diversification will also influence portfolio choice, and there will also be an interaction between taxation and risk, because taxes tend to dampen return fluctuations – after-tax returns have a lower variance than before-tax returns.

Taxation and Risk-Taking



To consider the effects of taxation on investment in risky assets, consider a model in which there are two states of nature (“good” and “bad”) and two assets, one safe (S), with return r^f in both states, and one risky (R), with returns $r^g > r^f$ and $r^b < r^f$ in the good and bad states. (Note that risk aversion requires that in equilibrium the expected return on the risky asset exceeds that on the safe asset.) The two assets together define the budget line as shown, and the tangency illustrated in the figure corresponds to a portfolio with positive holdings of both assets.


If income taxes at a single rate are imposed on all returns, this will contract the points R and S toward the point $(1,1)$ on the 45° line and lead to a parallel shift in the budget line and the new points R' and S' as shown, with relative asset demands influenced by an income effect. However, taxes may also affect the slope of the budget line; assuming that the risky asset is equity and the safe asset is debt, lower individual taxes on equity income would favor the risky asset, but capitalization of the generally more favorable equity taxation into lower before-tax equity returns would favor debt. Which effect dominates depends on whether the individual is a low-bracket or a high-bracket taxpayer (i.e., has a value of θ below or above θ^* in the first figure above). As a consequence, a substitution effect will push higher-bracket taxpayers toward holding more of the risky asset, and will push lower-bracket taxpayers toward holding less, but diversified portfolios will still be in order. How much of each asset is held will also be affected by risk aversion. For example, an infinitely risk-averse investor would have kinked indifference curves and would hold only the safe asset. A risk-neutral investor, with straight-line indifference curves, would choose to hold only the asset with the higher expected after-tax return.

Two other important points are worth making at this point about taxation and risk-taking, both of which can be related to the above two-state figure:

- (1) Tax systems typically treat gains and losses asymmetrically. While positive income is taxed, negative income (i.e., losses) does not receive a full tax refund. This means that the before-tax return on the risky asset in the good state may be $(1-\theta)r^g$, while the before-tax return in the bad state may simply be r^b , if $r^b < 0$, as is the case in the above figure. This would cause point R to shift horizontally to the left under taxation, to a point directly below R' , and hence to steepen the budget line and discourage investment in the risky asset.
- (2) Taxing risky assets reduces both expected returns and risk; the former discourages investment in risky assets, while the latter makes them more attractive. In one special case, the latter effect must dominate. Suppose that the tax system does not tax capital income generally, but just the excess over the safe rate of return, $(r^g - r^f)$ or $(r^b - r^f)$. Also assume that the tax system is symmetric, so that the issue just discussed does not arise. Then in the above two-state figure, taxation does not move point S , and simply shifts point R along the original budget line toward point S . For example, if the tax rate is 50%, R will move half-way from its original position to S . This does not change the investor's budget line, but must increase the portfolio share held in the risky asset. That is, a tax on *excess returns* – returns to risk-taking in excess of the safe rate of return – reduces a risky asset's expected return but in a way that does not change the investor's options and that encourages risk-taking. (A corollary is that the *expected* tax payment, which is positive, imposes no burden.) This encourages private risk-taking, but also increases the risk borne by the government, unless the government can pool the risks of individual tax payments; the transfer of that risk back to individuals may in turn reduce their risk-taking, and might even undo the initial increase in risk-taking, a point made by Gordon (*QJE* 1985) and Kaplow (*NTJ* 1994).

Evidence on the influence of taxes on portfolio choice is somewhat mixed but generally consistent with the theory that taxes should influence the mix of assets held. A relatively recent application is in the paper by Kawano, who studies the impact of the 2003 reduction in the US rate of dividend taxation, mentioned above. The effect was not only to make increase equity's

attractiveness relative to other assets, but also to increase the attraction of high-dividend-yield equities (having more income in the form of dividends rather than capital gains) relative to low-yield equities, especially for investors in higher tax brackets. Thus, we should have expected higher bracket investors, relative to low-bracket investors, to shift their portfolios toward high-dividend-yield stocks. In terms of the simple two-asset graph above, we can imagine the two assets being high-yield (h) and low yield stocks (l), with the tax rates on both stocks depending on the individual's ordinary tax rate, θ , but with high-yield stocks facing a higher tax rate ($h\theta > l\theta$), since dividends faced a higher tax rate than capital gains prior to 2003. (Even since 2003, other provisions make the effective capital gains tax rate somewhat lower than the dividend tax rate, but the gap is much smaller than before.) The 2003 legislation lowered the tax rate on both assets, but more for high-yield stocks, a benefit most valuable for those in high brackets.



As shown in the graph to the left, this change shifts the point of indifference to the right, given no change in before-tax returns. (We might also expect the before-tax return on high-yield stocks, r^h , to fall relative to that on low-yield stocks, r^l , in order to clear the market for the two types of stocks.) Indeed, Kawano finds a shift in portfolio sorting, with higher-bracket investors shifting more strongly toward high-yield stocks than low-bracket investors.

Capital Gains Taxation

Capital gains taxes are of particular interest for a number of reasons, even though they do not account for a large share of revenue for a typical government, including the United States. One reason for the interest is their concentration at the top of the income distribution. Another important aspect of capital gains is that they are taxed upon *realization* rather than on accrual, which makes the tax complex and subject to a variety of potential taxpayer responses.

What does realization-based taxation do? Consider a two-period model in which an investor has an asset purchased in an earlier period for \$1, which has already appreciated in value by an amount g . The investor can either hold the asset for another period, earning an additional return r , or sell and earn the market rate of return i . Suppose all income is taxed at rate t , but only when assets are sold. Also suppose that the investor's objective is to maximize terminal wealth.

If the investor sells the asset and reinvests, terminal wealth is:

$$W_R = (1+g(1-t))(1+i(1-t)) = (1+g)(1+i) - t[g(1+i(1-t)) + (1+g)i]$$

If the investor holds the asset until the end of the second period, terminal wealth is:

$$W_H = (1+g)(1+r) - t[(1+g)(1+r) - 1] = (1+g)(1+r) - t[g + (1+g)r]$$

Comparing the terms in brackets in the second version of each expression, we can see that the “hold” strategy enjoys a tax advantage over the “realize” strategy – first period gains, g , are taxed one period earlier under the latter, and hence the tax liability has a higher accumulated value at the end of the second period because it is multiplied by $1+i(1-t)$. It follows that if $i = r$, the investor will choose to hold rather than to realize, and indeed that there is a range of values of $r < i$ for which it will still be optimal to hold rather than to sell. This phenomenon is known as the lock-in effect – in order to defer tax on previously accumulated gains, individuals have an incentive not to sell assets even when, for non-tax reasons, they would prefer to sell. In this example, the lock-in effect is associated with the investor’s willingness to accept a lower before-tax rate of return, but in a realistic setting the major distortion comes from an inefficient allocation of assets across investors. That is, when an individual realizes a capital gain by disposing of an asset, that asset does not typically disappear, but instead ends up in someone else’s portfolio. Thus, it is unlikely simply to have a below-market rate of return, because asset prices adjust. Rather, in a setting with risky assets, other investors may be willing to pay more for the asset than the individual currently holding it. For example, suppose there are two investors with appreciated stock, one holding Apple and the other holding ExxonMobil. As returns on these two assets are not perfectly correlated, a combined portfolio would offer a better risk-return trade-off than either specialized position. Absent taxation, each investor could be made better off by trading with the other, but with a capital gains tax, the trades may not occur.

The lock-in effect is exacerbated by two other provisions found in the US tax system and typical of others as well. First, gains on assets held for at least one year are taxed at a lower rate (in United States at present, a maximum of 20% vs. a maximum tax rate on ordinary income of 37%). Second, gains held until death are not taxed at all. On the other hand, the lock-in effect is reversed when an asset has gone down in value ($g < 0$ in the above example), since deferral of tax in this case means deferring a tax *refund*. Thus, individuals have an incentive to hold gains and realize losses, meaning that those with large numbers of distinct positions in different assets could, on a regular basis, achieve liquidity by “harvesting” losses without having to realize gains. This possibility, in turn, is largely responsible for another tax provision, which limits the annual value of deductible losses (in excess of realized gains) to \$3,000. Unfortunately, a limit on the deductibility of losses also discourages risk-taking.

Empirical Evidence on Responses to Capital Gains Taxation

There has been a substantial literature relating capital gains realizations to capital gains tax rates. One of the key issues is the need to distinguish between short-run and long-run responses. We would expect that a change in tax rates could have a large impact on the timing of realizations, because individuals can adjust the timing of their asset sales. For example, after the October, 1986 passage of the Tax Reform Act of 1986, which increased the capital gains tax rate on high-income individuals from 20% to 28% effective January 1, 1987, there was such a surge in realizations in the remainder of 1986 that realizations for that year were approximately twice as high as those in 1985 or 1987. But that doesn’t mean that we would expect realizations to be permanently twice as high under a 20% tax rate as under a 28% tax rate.

One standard approach originally developed using panel data by Burman and Randolph (*AER* 1994) involves type-II Tobit estimation (for the decision to realize gains and gains realized), where the second, intensive-margin decision takes the form:

$$(1) \quad \ln g_{it} = \gamma_1(\tau_{it} - \tau_{it-1}) + \gamma_2\tau_{it}^p + \gamma_3(\tau_{it} - \tau_{it}^p) + X_{it}\gamma_4 + \varepsilon_{it}$$

where g is capital gains, X is a vector of individual attributes, τ is the individual's capital gains tax rate, and τ^p is a measure of the individual's "permanent" tax rate. The intuition for including the lagged tax rate τ_{it-1} is that a higher value will mean lower past realizations, hence a large stock of gains available to be realized at time t . The intuition for including some permanent tax rate measure, τ_{it}^p , is that if individuals expect a higher tax rate to prevail in the future, they will (as in 1986) wish to realize more gains in period t . But how should one represent this permanent tax rate? Following Auerbach and Siegel (*AER* 2000), Dowd et al. replace τ_{it}^p in the above specification with τ_{it+1} , the tax rate the individual will face the following year, which is generally known at time t , i.e., estimate the following expression

$$(1') \quad \ln g_{it} = \beta_1\tau_{it-1} + \beta_2\tau_{it} + \beta_3\tau_{it+1} + X_{it}\gamma_4 + \varepsilon_{it}$$

There is one further econometric issue that must be confronted in estimating (1'): the capital gains tax rate may depend on the level of gains realized, since tax rates rise with income. To deal with this, a common problem in empirical analysis of behavioral responses to taxation, all of these papers use as an instrument for τ a so-called "first-dollar" tax rate – the capital gains tax rate the individual would face on the first-dollar of capital gains realized. In their preferred specification, based on a panel of tax returns from 1999-2008, with tax rates incorporating both federal and state tax provisions, Dowd et al. find a permanent elasticity (corresponding to the effect $\beta_1 + \beta_2 + \beta_3$ in equation (1')) of -0.716 and a transitory elasticity (corresponding to the effect β_2) of -1.194, meaning that a temporary cut in the capital gains tax rate would increase tax revenue in the current year, but that a permanent cut in the capital gains tax rate would not. However, the impact of τ_{it+1} is insignificant – the larger response to a temporary tax cut comes through the impact of τ_{it-1} . This in contrast to the findings of Auerbach and Siegel as well as earlier time-series results, perhaps reflecting weaker anticipation effects in more recent years, which have lacked the dramatic policy changes of earlier periods.

Agersnap and Zidar (hereafter AZ) use state-level tax changes alone for their estimates. Their unit of observation is aggregate realizations by state and year, and their explanatory variable is the maximum capital gains tax rate by state and year, τ_{st} . This approach alleviates the need to control for sample selection (since aggregate realizations are always positive) or to correct for individual tax rate endogeneity (since the state maximum tax rate is arguably exogenous with respect to capital gains realizations). Using a more general specification than in expression (1') that allows for several leads and lags of the tax variable (and using the form $\log(1-\tau)$ rather than τ as the explanatory variable), AZ's estimates translate into a permanent elasticity with respect to τ of -0.53 – somewhat lower in magnitude than that found by Dowd et al., and again with no anticipation effects (i.e., there is no significant response to τ_{st+j} for $j>0$).

Reforming the Capital Gains Tax

Some changes in the capital gains tax (such as taxing capital gains at death) could serve to reduce the lock-in effect, but other problems remain as long as the basic approach to taxing capital gains upon realization is followed. Some arguments for keeping the capital gains tax rate lower than other capital income taxes, including the potentially higher behavioral response elasticity and the importance of capital gains in fostering venture capital investments, relate to

the realization-based nature of the tax (in the latter case because risky venture-capital investments face serious limitations on their ability to deduct losses, which as discussed earlier is a necessary feature of a realization-based system).

What other alternatives exist? One approach would be to tax capital gains at death, or at least to force heirs who receive assets to “carry over” the *basis* (i.e., original purchase price) of the assets received and therefore be liable for tax on the full gain when they eventually sell the assets. This would clearly reduce the lock-in effect associated with holding assets until death.

Another frequent proposal has been to index capital gains for inflation, allowing individuals to adjust an asset’s purchase price upward for changes in the price level since purchase (i.e., pay tax on the sale price V_t less the original purchase price, V_0 multiplied by the ratio of current and initial price levels, P_t/P_0). Letting π be the annual inflation rate, this would make the return to holding an asset, W_H , equal to:

$$W_H = (1+g)(1+r) - t[(1+g)(1+r) - (1+\pi)^2] = (1+g)(1+r) - t[(g - \pi)(1+\pi) + (1+\pi)(g - \pi)(r - \pi)]$$

I.e., real tax liability is independent of inflation for given real rates of return $(g - \pi)$ and $(r - \pi)$.

Perhaps the simplest idea for reform would be to tax capital gains as they accrue, rather than upon realization (perhaps combined with a reduced rate to offset the increased present value of taxes). But there are two problems with this approach: (1) taxpayers may lack liquidity to pay taxes until assets are actually sold; and (2) the government may not know the value of some assets until they are actually sold. One proposal for dealing with the liquidity problem, by Vickrey (*JPE*, 1939), amounts to keeping an account of accruing gains and the associated tax liability and charging interest on this accruing unpaid balance until asset sale. That is, the tax liability as of date s would evolve according to:

$$(2) \quad T_{s+1} = [1+i(1-t)]T_s + tr_sA_s$$

where r_s is the rate of return at date s , A_s is the value of the asset at date s , i is the safe rate of interest and t is the tax rate. A problem with Vickrey’s approach is that r_s and A_s may be unobservable, but Auerbach (*AER* 1991) argued that one can generalize Vickrey’s approach to:

$$(3) \quad T_{s+1} = [1+i(1-t)]T_s + tiA_s + t^*(r_s-i)A_s$$

where t^* can take on any value, since (as discussed above), a tax rate on a risky asset’s return in excess of the safe rate has no effect on the investor’s opportunities. Auerbach then showed that a tax liability of the form:

$$(4) \quad T_{s+1} = \left[1 - \left(\frac{1+i(1-t)}{1+i}\right)^s\right] A_s$$

satisfies (3). Note that only observable information is needed to assess the tax in (4): sale price, A_s , the holding period, s , the safe rate of interest, i , and the tax rate, t . Auerbach and Bradford generalize this result and show how it can be implemented using a tax system based exclusively on observed cash flows, without keeping track of individual assets and holding periods.