


Thus far our analysis of the effects of taxation has been largely within a static context. But there are many relevant issues, such as the role of expectations, the speed of adjustment, and the impact on different generations, that are difficult to consider without an explicit treatment of dynamics, i.e., how tax policy affects the economy over time.

Adding Dynamics to the Harberger Model

One question that arose in the analysis of the Harberger model was how one should think about the assumption that capital and labor fully adjust across production sectors in response to a tax change. Even if one maintains the assumptions of fixed factor supplies, full adjustment, particularly for capital, only makes sense in the long run. In the short run, it might make more sense to assume that labor adjusts but that capital does not. What would the implications be regarding incidence? If we impose a tax on corporate capital and capital does not initially move from the sector, it would seem that corporate capital, being temporarily immobile, bears the whole tax in the short run, and that with gradual adjustment the burden is shifted over time to all capital (for cases in which capital bears 100% of the tax in the long run). As discussed in the Auerbach's survey, we can trace the process in the following graph.



Suppose, at time t_0 , there is an *unannounced* tax, τ , on income from corporate capital. (If the tax change were anticipated, adjustment would begin before t_0 .) Initially, this causes a drop in the after-tax returns to capital in the corporate sector by the same amount as the tax, as the marginal product of capital in both sectors remains at r_0 . Over time, however, as capital shifts into the other sector, the marginal product of capital there falls, and the marginal product of corporate capital rises, until their after-tax returns are equated at some long-run value, r_∞^N . How long the adjustment takes depends on the costs of adjustment.

These changes in the after-tax return to capital over time, however, do not fully capture the incidence of the corporate tax, in terms of *who* bears the tax. It does not make sense to say that owners of noncorporate capital gradually bear more of the tax burden as adjustment occurs, because once the tax has been imposed, capital market equilibrium requires that corporate and noncorporate assets yield the same after-tax market *return*, which is distinct from the after-tax marginal *product* of capital. That is, the value of corporate capital at date t , q_t , must be such that the rate of return per dollar, including the after-tax return and the capital gain, equals that of noncorporate capital. The solution for the equilibrium path of q and capital adjustment will generally be unique once we impose an initial condition that the corporate and noncorporate

capital stocks are initially fixed and a terminal condition that the relative value of corporate capital converges to 1. One can trace out this adjustment path using phase diagrams, as discussed (using a somewhat different model) by Fullerton and Metcalf, pp. 1840-44. The path will be one on which the value of corporate capital initially drops below 1, reflecting the fact that corporate capital initially and for some time has a lower after-tax marginal product. This initial drop in value must be large enough so that the present value of after-tax returns to corporate capital and noncorporate capital are the same, per dollar of capital. That is, the integral of the gap between r^N and $r^G - \tau$ in the above figure must be capitalized as a discount in the initial value of a unit of corporate capital. Thus, a portion of all future corporate taxes is borne by initial shareholders. The remainder, which shows up in the decline over time in r^N , can be said to be borne over time by owners of all capital, since all purchasers of corporate *and* noncorporate capital after t_0 receive this rate of return.

Lifetime Incidence and Generational Accounting

Very often, conclusions we draw about incidence of taxes may be misleading if they are based on annual calculations. For example, it is common to assess tax burdens of different individuals by looking at the taxes they bear relative to current income. But current income may not be a particularly good indicator of an individual's ability to pay, as the following examples illustrate.

1. Under the permanent income/life-cycle models of consumption behavior, individuals smooth consumption – consumption fluctuates less than income. This means that the consumption-income ratio will fall with income in any given year, even if consumption is a constant share of permanent or lifetime income. Thus, assessments using annual data will tend to overstate the regressivity of consumption taxes.
2. Like many old-age pension systems, the US social security system imposes payroll taxes during working years and pays benefits after retirement. As incomes fall in retirement, using annual income to assess ability to pay will make the system look very progressive, as it is taxing “high-income” workers to fund transfer payments to “low-income” retirees. But, on a lifetime basis, one's conclusions might be very different, as the retirees might have been as affluent while working as those being taxed to finance their benefits.

As the second illustration shows, a further complication arises when different generations are involved, because even if we use a longer-run measure of ability to pay, we still have a problem of assessing burdens when there may be transfers among generations. How can we say whether the social security system is progressive if the taxes and transfers within one generation do not balance? Clearly, we need to consider the distribution of tax burdens not only within generations, but also across them. This is what generational accounting endeavors to do.

Typically, we use accumulations of national debt as a shorthand indicator of the extent to which fiscal burdens are being transferred from current generations to future ones. But this is not a very accurate indicator, and the growing importance of age-based government policies (such as old-age pensions and medical care) further limits its usefulness.

Consider again the US social security system. This system is run largely on a pay-as-you-go (PAYG) basis, meaning that today's taxes go to pay today's benefits; even though it is often

described as a contributory pension scheme, individuals are not funding their own future benefits. A trust fund has been accumulated through the years, standing (according to the August, 2021 Social Security Trustees Report, Table II.B1) at \$2.9 trillion at the end of 2020, but it is small relative to the system's remaining, unfunded liability (equal to the present value of benefits less taxes – the so-called open group liability of the system – and less trust fund assets) under current rules, which was \$59.8 trillion (Table VI.F2), up from \$53.0 trillion one year earlier – a net annual deficit of \$6.8 trillion! This compares to the US government's overall official budget deficit of \$3.1 trillion during 2020, which itself was extraordinarily high due to Covid-19. Yet, the budget of the social security system showed a small budget *surplus* of \$10.9 billion, not a deficit of \$6.8 trillion, because the trust fund increased slightly over the course of the year; the increase in expected future benefits net of taxes is ignored. As first pointed out by Feldstein (*JPE* 1974), this implicit liability is like national debt in another important respect; we would expect individuals to perceive the right to receive social security benefits as an addition to wealth, just as ownership of government bonds would. (In each case, the wealth effect presumes that individuals do not view future taxes on subsequent generations as if they were taxes on themselves, as they would under Ricardian equivalence.) Note that if the social security system were run differently, for example if individuals were issued government bonds in exchange for their payroll taxes and could redeem the bonds to provide an income flow during their retirement, the implicit liability would be converted into an explicit one.

The construction of generational accounts is intended to overcome the ambiguity of government debt as a measure of intergenerational transfers. We start with the identity relating government debt at the beginning of period t and the components of annual deficits, government purchases, G_t , taxes net of transfer payments, T_t , and interest on the national debt, rD_t (where for simplicity we assume that r is constant over time):

$$(1) \quad D_{t+1} = G_t - T_t + (1 + r)D_t$$

Solving this difference equation forward and imposing the terminal condition that the government cannot run a Ponzi game (that is, $(1+r)^{-T}D_{t+T} \rightarrow 0$ as $T \rightarrow \infty$), we get the government intertemporal budget constraint (GIBC):

$$(2) \quad \sum_{s=t}^{\infty} (1+r)^{-(s-t+1)} T_s = D_t + \sum_{s=t}^{\infty} (1+r)^{-(s-t+1)} G_s$$

Now, break the components of T_t at each date into values for each cohort alive at that time,

$$(3) \quad T_t = \sum_{k=t-D}^t T_t^k$$

where k indexes the cohort's year of birth and D is lifespan. Finally for each cohort, k , take the present value of these annual terms, from either the current year or the cohort's year of birth, whichever is later, to form that cohort's generational account:

$$(4) \quad N_{t,k} = \sum_{j=t}^{k+D} (1+r)^{-(j-t+1)} T_j^k \quad \forall k \leq t; \quad N_{k,k} = \sum_{j=k}^{k+D} (1+r)^{-(j-k+1)} T_j^k \quad \forall k > t$$

Note that the terms $N_{t,k}$ and $N_{k,k}$ in (4) account for all components of taxes from date t forward, so we can rewrite the GIBC:

$$(5) \quad \sum_{k=t-D}^t N_{t,k} + \sum_{k=t+1}^{\infty} (1+r)^{-(k-t)} N_{k,k} = D_t + \sum_{s=t}^{\infty} (1+r)^{-(s-t+1)} G_s$$

(Here, we've assumed that government purchases are not allocated to generational accounts, but an alternative would be to allocate at least some components of G as well.)

Returning to the issue of implicit liabilities, note that if we changed the accounting for social security, treating payroll taxes and purchases of government bonds and benefits as receipts of interest and principal on these bonds, then the value of D_t would increase, the values of $N_{t,k}$ for current generations would decrease by the same amount in present value, but the generational accounts for future generations would be unaffected.

We can measure the government's fiscal imbalance by assuming that current policy is maintained for all existing generations and asking by what fraction the generational accounts of future generations would have to be inflated, relative to current policy, to ensure that the equality in (5) is satisfied. Note that this calculation would not be affected, for example, by a change in accounting convention that converted implicit liabilities to explicit ones.

Application: Annual and Lifetime Inequality

The degree of inequality, and the extent to which government taxes and transfers contribute to or mitigate this inequality, are questions of major importance, especially in the United States, where inequality has risen in recent decades. But, as the above discussion of the social security system illustrates, grouping together individuals at different points in the life cycle can lead to misleading answers both about the extent of inequality (e.g., as we will be treating retirees with low current income as poor) and the effects of government policy on inequality (e.g., we will be treating social security benefits as payments to the poor). Auerbach, Kotlikoff, and Koehler deal with both issues by estimating generational accounts, on a remaining lifetime basis, for individuals in different age cohorts at different places in the lifetime resource distribution, before taxes and transfers (where resources equal current wealth plus the present value of projected future labor income). They find that the degree of progressivity of the fiscal system is understated by looking at current-year, rather than lifetime, taxes and transfers.

Capital Income Taxes, Labor Income Taxes and Consumption Taxes

When thinking about the optimal taxation of saving and capital income, dynamic issues naturally arise. Of course, we could start by simply reinterpreting our existing optimal tax results by letting different consumption goods represent consumption at different dates. That is, for a representative individual, the three-good model with two consumption goods and leisure might be thought of as a model in which the individual chooses how much to work when young, how much of labor earnings to consume immediately, and how much to save for retirement consumption. If the individual earns wL when young, consumes C_1 and saves $(wL - C_1)$, that individual's second period consumption will be $(wL - C_1)(1+r)$, where r is the one-period rate of return. We may rewrite this budget constraint (letting labor be numeraire) as:

$$(1) \quad \frac{1}{w}C_1 + \frac{1}{w(1+r)}C_2 = L,$$

from which it is obvious that a tax on capital income (by reducing r after-tax) will impose a higher tax rate on second-period consumption than first-period consumption. Thus, taxing capital income at a positive rate would be called for only if second-period consumption is more

complementary to leisure than first-period consumption. Under equal complementarity, we prefer simply to tax labor income or, equivalently, to impose a uniform tax on consumption.

We could extend this to a case in which individuals differ in both ability and preferences, in which case distributional concerns might dictate a higher tax on second-period consumption if it tends to be concentrated more among higher-ability individuals. But there are many things missing from this translation of static analysis that can have a significant impact on our conclusions. Here is a non-exhaustive list:

Multiple Periods

Adding several periods of saving and consumption leads to stronger conclusions about the desirability of capital income taxation. This might seem surprising, since the standard optimal tax model's results don't really change when we move from two consumption goods to several, but the particular way in which capital income taxes translate into consumption price distortions is what matters here. Suppose we extend expression (1) to cover several periods of consumption, still with just one period of labor. The budget constraint then becomes:

$$(1') \quad \frac{1}{w}C_1 + \frac{1}{w(1+r)}C_2 + \dots + \frac{1}{w(1+r)^{T-1}}C_T = L,$$

from which it is clear that a constant rate of capital income tax will distort the prices of future consumption more and more as T grows. Even if we wish to tax future consumption at a higher rate than current consumption, the optimal differential tax rate is unlikely to grow without bound as T increases. Thus, at some point in the future the capital income tax has to converge to zero, to prevent the distortion from continuing to grow. This is the intuition of the Chamley (*Econometrica* 1986)-Judd (*Journal of Public Economics* 1985) result that with an infinite horizon consumer the optimal capital income tax converges to zero in the long run.

Bequests

If all bequests are altruistic, as in the Barro-Becker view, then having a short life span with an operative bequest motive is just like having a single individual with a long planning horizon, which is the case just considered. But this logic only holds if bequests are positive. Also, there may be other motivations for leaving bequests. They can arise from a joy of giving motivation, as an accident of saving for retirement when lifetimes are uncertain and annuities markets imperfect, or from a strategic motive involving intrafamily bargaining. These have different implications for capital income taxation, either directly or indirectly through the taxation of bequests. Also, to the extent that the welfare of individuals who receive bequests receipts is not fully internalized by bequest motives, there may be externalities associated with bequests.

Dynamic Inconsistency

Unlike in the Arrow-Debreu world, where decisions are made once, even with respect to purchases of goods at future dates, governments can deviate from an announced policy once individuals have made decisions that reduce their options. This is the standard problem of dynamic inconsistency, and can lead government to increase capital income taxes *ex post* above their optimal *ex ante* levels. For example, in the Chamley-Judd set-up, it is optimal to tax capital income very heavily in the short run, when capital is essentially in fixed supply, and then have the capital income tax rate fall over time. But this will also be the optimal policy, starting from any future date, meaning that the initial policy of low future tax rates is dynamically inconsistent

unless there is some mechanism under which government can limit its future ability to deviate from previously announced policy.

Liquidity Constraints

As with dynamic inconsistency, this is a factor that would not arise in the Arrow-Debreu world; liquidity constraints impose additional constraints beyond an overall lifetime budget constraint on the allocation of lifetime resources over consumption at different dates, because once money has been borrowed against future income, individuals have an incentive to change their plans for future dates in a way that jeopardizes repayment. If liquidity constraints apply, perhaps to individuals early in the life cycle, then taxes that would otherwise be equivalent no longer will be. For example, the equality between equal-present-value labor income taxes and consumption taxes breaks down, because labor income taxes are paid earlier in life and therefore may exacerbate the restrictions imposed by liquidity constraints.


Transitions

Tax policy takes place in real time. This means that expectations matter (as already mentioned in the context of analyzing incidence in the Harberger model). It also means that individuals at different ages will be affected differently, a fact that analysis of how a representative individual or cohort is affected by a tax change won't capture. Taking transitions into account can alter our conclusions about the desirability of different tax systems.

Consider again the two-period budget constraint, but this time suppose that the representative individual has initial assets,

$$(6'') \quad C_1 + \frac{1}{(1+r)} C_2 = wL + A$$

It is evident that a consumption tax base would be broader than that of a labor income tax, hitting consumption from existing assets. This initial endowment of wealth could be an inheritance, but in a transition that hits some generations during their lifetimes, it could also simply represent previous own accumulations of life-cycle wealth. Thus, a consumption tax imposed at any date will have a broader tax base, in present value, than a labor income tax, the difference equal to existing wealth in the economy. This can have a big impact on outcomes, as shown by Auerbach and Kotlikoff (*Dynamic Fiscal Policy* 1987). Equivalent variations (relative to lifetime resources) have the following pattern, for an immediate switch from a 15 percent income tax to either a consumption tax or a labor income tax, maintaining the same annual revenue levels (with generation 0 being the one reaching adulthood in the year the transition begins):



The taxes have different effects on older initial generations, as labor income taxes reduce their taxes while consumption taxes increase them. This difference translates into differences among future generations, who face lower taxes as a result under consumption taxes and higher taxes under labor income taxes. We can neutralize these differences using a hypothetical system of balanced-budget lump-sum taxes and transfers among generations, so that transition generations are held harmless and future generations share equally in gains or losses, but this eliminates only some of the difference. The remainder is due to the fact that taxing initial wealth not only has effects on generational incidence, but also on efficiency – it's a lump-sum tax, if not anticipated. In summary, both consumption taxes and labor income taxes reduce intertemporal distortions, but for labor income taxes the windfall transfers to initial generations involve an efficiency loss that more than offsets the efficiency gains associated with a reduced intertemporal distortion. Altig et al. provide an extension for more realistic tax systems and within-cohort heterogeneity.

Many Periods of Labor Supply

Even if we do not wish to tax consumption at different dates differently, and therefore may wish to tax capital income at a low rate, we might still wish to tax labor income differently at different ages, for example if labor supply elasticities differ between prime-age workers and those just entering or leaving the labor force. If there are constraints on age-based labor income taxation, then it is possible that we would wish to use capital income taxation as a proxy. For example, since consumption from wealth is more important for older individuals, we might proxy for higher labor income taxation of middle-aged individuals by taxing their subsequent consumption effectively through capital income taxation. See Erosa and Gervais (*JET* 2002), who find that capital income taxes should be positive for this reason, but still much lower than labor income taxes.

Uncertain Earnings

This also relates to taxing earnings at different dates at different rates. Suppose lifetime earnings are uncertain; some individuals have more favorable draws than others in terms of their labor market outcomes. We might wish to tax future labor income at higher marginal rates, to provide insurance (which taxing labor income at younger ages would not do), but if we cannot do so we might again wish to tax capital income to simulate higher labor income tax rates among those who will have future consumption financed by saving. See Conesa, Kitao and Krueger (*AER* 2009).

Behavioral Issues

Retirement saving is one element of behavior that is often thought to be subject to deviations from the standard rational choice model. Individuals must make decisions using a very long horizon, and they do not get to learn from their own mistakes. Therefore, some policies that might seem suboptimal, such as tax incentives to place money into retirement accounts from which withdrawals are restricted, could become desirable.

New Dynamic Public Finance

Although we've deferred until 230B the general optimal income tax problem, we note here that its basic approach is to choose a marginal rate schedule subject to self-selection constraints. That is, with no other constraints on the shape of the marginal tax rate schedule, we can raise taxes on higher income individuals up to the point where any further increases would cause them to prefer to represent themselves as having lower skills. The basic Mirrlees (*REStud* 1971) model is a static one, in which there is no capital income. With more than one period, what role should capital income taxation play? The central insight of the NDPF is that capital income taxation may aid in the imposition of progressive labor income taxation, by discouraging individual saving. The intuition is that high-skill individuals who have saved in previous periods will find it easier not to work, for they can rely on wealth (and the government's more favorable treatment of low-income individuals) to finance a reasonable level of consumption. With limited saving, the costs of reducing their labor supply will be greater and therefore higher labor income tax rates can be imposed on them. See the discussion of this and related results in the review by Stantcheva.