Volume Title: Generational Accounting around the World

Volume Author/Editor: Alan J. Auerbach, Laurence J. Kotlikoff and Willi Leibfritz, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-03213-2

Volume URL: http://www.nber.org/books/auer99-1

Publication Date: January 1999

Chapter Title: The Methodology of Generational Accounting

Chapter Author: Alan J. Auerbach, Laurence J. Kotlikoff, Willi Leibfritz

Chapter URL: http://www.nber.org/chapters/c6684

Chapter pages in book: (p. 31 - 42)

2 The Methodology of Generational Accounting

Alan J. Auerbach and Laurence J. Kotlikoff

This chapter describes the standard method of generational accounting that is used, with minor modifications, in all the country studies. This methodology was first developed in Auerbach, Gokhale, and Kotlikoff (1991), on which this chapter closely draws.

Generational accounting is based on the government's intertemporal budget constraint. This constraint, written in equation (1), requires that the future net tax payments of current and future generations be sufficient, in present value, to cover the present value of future government consumption as well as service the government's initial net indebtedness.[1]

$$(1) \qquad \sum_{k=t-D}^{t} N_{t,k} + (1+r)^{-(k-t)} \sum_{k=t+1}^{\infty} N_{t,k} = \sum_{s=t}^{\infty} G_s (1+r)^{-(s-t)} - W_t^g .$$

The first summation on the left-hand side of equation (1) adds together the *generational accounts*—the present value of the remaining lifetime net payments—of existing generations. The term $N_{t,k}$ stands for the account of the generation born in year k. The index k in this summation runs from $t - D$ (those aged D, the maximum length of life, in year 0) to t (those born in year 0).

The second summation on the left-hand side of equation (1) adds together the present values of the generational accounts of future generations, with k again representing the year of birth. As each of these generational accounts is expressed in dollars of the respective generation's birth year, they must be

Alan J. Auerbach is the Robert D. Burch Professor of Economics and Law at the University of California, Berkeley, and a research associate of the National Bureau of Economic Research. Laurence J. Kotlikoff is professor of economics at Boston University and a research associate of the National Bureau of Economic Research.

1. The constraint does not assume that government debt is ever fully paid off, merely that the debt grows less quickly than the rate of discount—that it does not explode. Thus it is consistent with the long-run existence of government deficits, as long as these deficits are smaller than the amount needed simply to service the level of outstanding debt.

discounted back to year t in the summation, using the government's real before-tax return r.

The first term on the right-hand side of equation (1) expresses the present value of government consumption. In this summation the values of government consumption in year s, given by G_s, are also discounted to year t. The remaining term on the right-hand side, W_t^g, denotes the government's net wealth in year t—its assets minus its explicit debt.

Equation (1) indicates the zero-sum nature of intergenerational fiscal policy. Holding the present value of government consumption fixed, a reduction in the present value of net taxes extracted from current generations (a decline in the first summation on the left-hand side of eq. [1]) necessitates an increase in the present value of net tax payments of future generations.

The generational account $N_{t,k}$ is defined by

$$N_{t,k} = \sum_{s=\kappa}^{k+D} T_{s,k} P_{s,k} (1 + r)^{-(s-\kappa)}, \tag{2}$$

where $\kappa = \max(t,k)$. In expression (2) $T_{s,k}$ stands for the projected average net tax payment to the government made in year s by a member of the generation born in year k. The term $P_{s,k}$ stands for the number of surviving members of the cohort in year s who were born in year k.[2] For generations who are born prior to year t, the summation begins in year t and is discounted to year t. For generations who are born in year $k > t$, the summation begins in year k and is discounted to that year.

A set of generational accounts is simply a set of values of $N_{t,k}$, one for each existing and future generation, with the property that the combined present value adds up to the right-hand side of equation (1). Though we distinguish male and female cohorts in the results presented below, we suppress sex subscripts in equations (1) and (2) to limit notation.

Note that generational accounts reflect only taxes paid less transfers received. With the exception of government expenditures on health care and education, which are treated as transfer payments, the accounts do not impute to particular generations the value of the government's purchases of goods and services because it is difficult to attribute the benefits of such purchases. Therefore, the accounts do not show the full net benefit or burden that any generation receives from government policy as a whole, although they can show a generation's net benefit or burden from a particular policy change that affects only taxes and transfers. Thus generational accounting tells us which generations will pay for government spending not included in the accounts, rather than

2. As discussed in chap. 6 in this volume, by Ablett, the population weights $P_{s,k}$ incorporate both mortality and immigration, implicitly treating immigration as if it were a "rebirth" and assigning the taxes paid by immigrants to the representative members of their respective cohorts. This approach does not, therefore, separate the burdens of natives and immigrants. Such an extension is desirable, particularly if one wishes to study the effects on generational accounts of changes in immigration patterns.

telling us which generations will benefit from that spending. This implies nothing about the value of government spending; that is, there is no assumption, explicit or implicit, concerning the value to households of government purchases.

2.1 Assessing the Fiscal Burden Facing Future Generations

Given the right-hand side of equation (1) and the first term on the left-hand side of equation (1), we determine, as a residual, the value of the second term on the left-hand side of equation (1), which is the collective (aggregate) payment, measured as a time t present value, required of future generations. Based on this amount, we determine the average per capita present value lifetime net tax payment facing members of each future generation under the assumption that the average per capita lifetime net tax payment of members of successive generations rises at the economy's rate of productivity growth. This makes the lifetime payment a constant share of lifetime income. Controlling for this growth adjustment, the average per capita lifetime net tax payments of future generations are directly comparable with those of current newborns, since the per capita generational accounts of both newborns and individual members of future generations take into account net tax payments over these agents' entire lifetimes and are discounted back to their respective years of birth.

Our assumption that the generational accounts of all future generations are equal, on a per capita basis, except for a growth adjustment, is just one of many assumptions we could make about the distribution across future generations of their collective net payment to the government. We could, for example, assume a phase-in of the additional fiscal burden (positive or negative) to be imposed on future generations, allocating a greater share of the burden to later future generations and a smaller share to earlier ones. Clearly, such a phase-in would mean that generations born after the phase-in period has elapsed would face, on a per capita basis, larger values of lifetime burdens (the $N_{t,k}$) than we are calculating here.

Another way of measuring the imbalance of fiscal policy, illustrated in the chapters that follow, is to ask what permanent change in some tax or transfer instrument, such as an increase in income taxes or a reduction in old-age social security benefits, would be necessary to equalize the lifetime growth-adjusted per capita fiscal burden facing current newborns and future generations. Because such policies satisfy the government's intertemporal budget constraint, they are also sustainable.

2.2 Assumptions Underlying Generational Account Calculations

To produce generational accounts, we require projections of population, taxes, transfers, and government expenditures; an initial value of government wealth; and a discount rate. We consider the impact of total, not just national, government.

Typically, we assume that government purchases grow at the same rate as GDP, although in some cases we break these purchases down into age-specific components and assume that each component remains constant per member of the relevant population, adjusted for the overall growth of GDP per capita. This causes different components of government purchases to grow more or less rapidly than GDP according to whether the relevant population grows or shrinks as a share of the overall population.

Government infrastructure purchases are treated like other forms of purchases in the calculations. Although such purchases provide an ongoing stream rather than a one-time amount of services, they must still be paid for. Generational accounting clarifies which generation or generations will have to bear the burden of these and other purchases. For government wealth, we measure the government's net financial assets—its financial assets less its gross debt. We do not include the real assets of state enterprises in this measure but instead subtract projected net profits from state enterprises from projected government spending. This procedure effectively capitalizes the value of these enterprises.

Government wealth does not include the value of the government's existing infrastructure, such as parks. Including such assets would have no impact on the estimated fiscal burden facing future generations because including these assets would require adding to the projected flow of government purchases an offsetting flow of imputed rent on the government's existing infrastructure.

Taxes and transfer payments are each broken down into several categories. Our general rule regarding tax incidence is to assume that taxes are borne by those paying the taxes, when the taxes are paid: income taxes on income, consumption taxes on consumers, and property taxes on property owners. There are two exceptions here, both of which involve capital income taxes. First, we distinguish between marginal and inframarginal capital income taxes. As described below, inframarginal capital income taxes are distributed to existing wealth holders, whereas marginal capital income taxes are based on future projected wealth holdings. Second, in the case of small open economies, marginal corporate income taxes are assumed to be borne by (and are therefore allocated to) labor. The need for this later adjustment is discussed in chapter 15, for example.

The typical method used to project the average values of particular taxes and transfer payments by age and sex starts with government forecasts of the aggregate amounts of each type of tax (e.g., payroll) and transfer payment (e.g., welfare benefits) in future years. These aggregate amounts are then distributed by age and sex based on cross-sectional relative age-tax and age-transfer profiles derived from cross-sectional microdata sets. For years beyond those for which government forecasts are available, age- and sex-specific average tax and transfer amounts are assumed to equal those for the latest year for which forecasts are available, with an adjustment for growth.

2.3 Calculating Inframarginal Capital Income Taxes

Capital income taxes require special treatment because, unlike other taxes, they may be capitalized into the values of existing assets. Also, the time pattern of income and tax payments may differ. As a result of these features of capital income taxes, such taxes must be attributed with care to ensure that they are assigned to the proper generation. If all forms of capital income were taxed at the same rate, there would be no such problem: all assets would yield the same rate of return before tax (adjusted for risk), and each individual would face a rate of return reduced by the full extent of the tax. However, if tax rates on the income from some assets, typically older ones, are higher than those facing income from new assets (e.g., because of investment incentives targeted toward new investment), a simple arbitrage argument indicates that the extra tax burden on the old assets should be capitalized into these assets' values.

To illustrate the nature of the necessary correction, consider the case of cash-flow taxation in which assets are written off immediately. A well-known result is that the effective marginal capital income tax rate under cash-flow taxation is zero. However, taxes would be collected each year on existing capital assets, and such assets should therefore be valued at a discount. Assigning these taxes to the assets' initial owners, rather than to members of future generations who may purchase the assets, is consistent with the fact that such future generations of individuals may freely invest in new assets and pay a zero rate of tax on the resulting income. Our correction to actual tax payments should, in this case, result in a zero tax burden on the income from new assets.

For the general case, we use the following methodology. Our calculation begins with expression (3) for the user cost of capital, to which firms set their marginal products:

$$C = \frac{(r + \delta)(1 - k - \tau z)}{(1 - \tau)}, \tag{3}$$

where r is the investor's required after-tax return, δ is the investment's economic rate of depreciation, τ is the investor's marginal tax rate, k is the investment tax credit or grant received on investment, and z is the present value of depreciation allowances. We wish to calculate two measures. The first, which we denote by Q, is the tax-based discount on old capital, which equals the difference between tax savings from depreciation allowances and investment credits per unit of new capital and those available per unit of existing capital:

$$Q = k + \tau(z - z^{o}), \tag{4}$$

where z^{o} is the present value of depreciation allowances per unit of old capital.

Measured capital income tax payments are not based on the effective rate of tax on new capital m, where

(5) $$m = \frac{C - (r + \delta)}{C - \delta}.$$

Instead, they are based on an average tax rate, α, where

(6) $$\alpha = \frac{\tau(C - b) - k}{C - \delta}$$

and b is the average current depreciation deduction per unit of total capital. Comparing equations (5) and (6) indicates that we must correct measured taxes per unit of capital by subtracting from $\alpha(C - \delta)$ the term Δ, where

(7) $$\Delta = (\alpha - m)(C - \delta).$$

The values of z^o and b depend on past patterns of investment and the depreciation schedules permitted existing assets. For the case in which investment grows smoothly at rate n and all capital (new and old) is written off at rate ψ based on historic asset cost, the value of undepreciated basis per unit of existing capital may be shown to equal

(8) $$\frac{n + \delta}{n + \pi + \psi},$$

where π is the rate of inflation. Thus the value of b, the average current depreciation deduction per unit of capital, is ψ multiplied by this basis:

(9) $$b = \psi \frac{n + \delta}{n + \pi + \psi},$$

and the value of z^o, the present value of depreciation deductions per unit of existing capital, equals

(10) $$z^o = z \frac{n + \delta}{n + \pi + \psi},$$

where z is the present value of depreciation deductions per unit of basis (and per unit of new capital),

(11) $$z = \frac{\psi}{r + \pi + \psi}.$$

Substituting equations (5), (6), (9), and (11) into equation (7), we obtain

(12) $$\Delta = (r + \delta)\tau z \left[1 - \frac{(r + \pi + \psi)(n + \delta)}{(n + \pi + \psi)(r + \delta)}\right].$$

Substituting equation (10) into equation (4), we obtain

(13) $$Q = k + \tau z \left(1 - \frac{n + \delta}{n + \pi + \psi}\right).$$

Based on parameter values for the United States in the 1990s, Auerbach et al. (1991) estimated values of $\Delta = .00111$ and $Q = .111$.

There are other possible assumptions we could make about the incidence of capital income taxes. For a small open economy, for example, it may make sense to assume that taxes on mobile corporate capital are borne by local, fixed factors such as labor.[3]

2.4 Discount Rates and Uncertainty

For base-case calculations, generational accounts typically use a real rate of discount in the neighborhood of 5 percent, a rate that exceeds the real government short-term borrowing rate in most developed countries. This rate seems justified given the riskiness of the flows being discounted. However, as we now discuss, the "right" discount rate to use is in sufficient question to merit presenting results based on a range of alternative discount rates—a practice routinely followed by generational accountants.

The appropriate discount rate for calculating the present value of future government revenues and expenditures depends on their uncertainty. If all such flows were certain and riskless, it would clearly be appropriate to discount them using the prevailing term structure of risk-free interest rates. However, even in this simple and unrealistic case, such discounting could be problematic since it would require knowing the values of this term structure. To discern these values, one might examine the real yields paid on short-term, medium-term, and long-term inflation-indexed government bonds. But this presupposes the existence of such bonds. Many countries do not issue indexed bonds, and those that do don't necessarily issue indexed bonds of all maturities. The United States is a case in point. It has just begun to issue indexed bonds but so far has limited its issue of such bonds to those with 10-year maturities. Even if a country issues indexed bonds of multiple maturities, equating their real yields with the risk-free rate requires assuming no default risk, which for many countries is a very strong assumption.

In the realistic case in which countries' tax revenues and expenditures are uncertain, discerning the correct discount rate is even more difficult. In this case, discounting based on the term structure of risk-free rates (even if it is observable) is no longer theoretically justified. Instead, the appropriate discount rates would be those that adjust for the riskiness of the stream in question. Since the riskiness of taxes, spending, and transfer payments presumably differ, the theoretically appropriate risk-adjusted rates at which to discount taxes, spending, and transfer payments would also differ.

Is risk adjusting really important? A priori, one might think that forming the expected present value of future taxes and transfers of current and future generations, with discounting done at risk-free rates, would yield a meaningful

3. This approach is taken in Auerbach et al. (1997) for New Zealand.

measure of the fiscal burdens facing different generations on average.[4] But this is not the case as the following line of argument, relying on the invariance of economic outcomes under changes in fiscal labels, makes clear.

Chapter 1 points out that there are an infinite number of ways to label a country's underlying fiscal policy. If economic agents are rational, the choice of labels will have no real impact, including no impact on the intergenerational distribution of well-being—which generational accounting seeks to help illuminate. This proposition that economic outcomes are invariant under changes in the government's vocabulary is true regardless of whether the economy features uncertainty, including uncertain government policy. However, in the context of uncertain government policy, relabeling fiscal policy can easily alter expected future taxes and expected future transfer payments. Such relabeling will also alter the riskiness of reported taxes and transfer payments and, therefore, the proper rates at which to discount expected future tax and transfer streams. If one discounts these altered expected values with the proper risk-adjusted discount rates, one finds what one should find: no change in the expected utilities of any generation. However, if one simply uses the time path of risk-free rates of return to discount the expected value of future taxes and transfers, one gets nonsensical results: the "expected" fiscal burdens facing alternative generations depend on how the fiscal policy is labeled; that is, they depend on the government's choice of vocabulary.

An example may help clarify this point. Take the case of a fully funded defined-benefit social security system. Suppose the government has held risk-free bonds and now chooses instead to invest in risky stock. To acquire the stock, the government sells the public its bonds. Consequently, the public ends up holding stock through the government and bonds in its private portfolio. If the stocks perform well, the government rebates to the public (in the form of a transfer payment) the amount beyond what is needed to cover its social security pension obligations. If the stocks perform poorly, the government taxes the public to cover its social security obligations.

Hence, under the "new" policy, the public receives a sure income on the bonds that it has purchased from the government but a risky stream associated with the transfers or taxes it now faces. On balance, the public ends up with exactly the same income; that is, it gets the same social security pension income, and the combination of its safe bond income and its now risky net taxes is equivalent to its holding directly the stock sold to the government. This "portfolio" change on the part of the government alters the expected net tax payments of the public but has no real effects—it is nothing more than a relabeling of government receipts and payments. The fact that the government and

4. Diamond's (1996) presentation of "projections" seems to come close to endorsing such analysis, although Diamond's main argument is the same as we make here, namely, that properly valuing uncertain tax and transfer flows requires adjusting for risk when one discounts.

the private sector exchange different securities with the public is simply part of the relabeling process, not evidence that policy has fundamentally changed.

Another issue that arises with respect to risk-adjusted discounting is that the proper risk adjustments may be generation specific. To see this, consider a two-period model in which there are two generations and no government purchases—just an initial stock of debt that needs to be serviced and repaid in the second period. Generation 1, currently alive, will pay some tax rate, τ, times its uncertain income, and generation 2, not yet alive, will pay the residual. Since, by construction, the payments of the two generations equal principal plus interest on the debt in every state of nature, the government's intertemporal budget constraint is always satisfied.

In this example, aggregate tax payments are certain, although each generation's own tax payments are not. For generation 1, the uncertainty of its tax payments are actually a plus, since its risky income is being insured. Thus we would be justified in applying a discount rate of $\rho > r$, where r is the risk-free rate, in valuing the expected tax payments from generation 1's perspective. From generation 2's point of view, the situation is more complicated. It depends on how much generation 2's marginal utility of consumption is correlated with that of generation 1. If there were perfect correlation (say, because of a single source of income or complete intergenerational risk sharing), then generation 2's burden would be *greater* than that implied by discounting at the risk-free rate—its burden would be relatively higher in bad (low income) states of nature—so its expected tax payments should be discounted at a rate $\rho < r$.

Hence, by discounting the burdens of each generation at an appropriate discount rate (higher than the risk-free rate for generation 1, lower than the risk-free rate for generation 2), we would still find that the sum of the burdens satisfied the government's intertemporal budget constraint but get a better measure of the impact on individual utility.

To see the implications of this result, let us go back to the general, multiperiod and multigeneration model and assume again for the moment that there is a single set of state-contingent future prices that all generations use to evaluate future flows. Then our current approach, to define the burden on future generations as a residual, gives a correct measure of the aggregate burden on future generations. That is, we define this collective burden as

$$(18) \qquad N_{\text{fut}} = \sum_{k=-D}^{0} N_{0,k} - \sum_{t=0}^{\infty} (1 + \rho)^{-t} G_t - B_0,$$

where $N_{0,k}$ is the generational account for the generation born in year k, formed by discounting that generation's flows using the discount rate ρ. As long as ρ is chosen appropriately (as already discussed, this would include, perhaps, using different values of ρ for government purchases than for taxes and transfers), N_{fut} is the value of the burden placed on future generations. Note that this procedure *will not* give us a measure of the expected values of net tax payments by future generations, but rather the value of these payments based on the valuation that would be placed on such payments by existing generations.

But now we must come back to consider how to value the residual flows that must be paid by future generations under incomplete risk sharing across generations. Consider again the simple model with two agents. In this instance, we cannot use a discount rate based on generation 1's valuation of generation 2's burden. If we evaluate generation 2's burden from its own point of view, the burden may be lower. For example, suppose that the income of the two generations is negatively correlated; the negative correlation might arise if, for example, the source of shocks was to the relative productivity of capital and labor and generation 1 (2) supplied capital (labor). Then generation 2's burden, from its own perspective, will be less onerous than a certain burden with the same expected value. This is because generation 2's taxes will be higher in states in which its income is higher (even though generation 1's income is lower). Hence, we should discount generation 2's expected burden at a rate higher than the risk-free rate. Thus both generations will perceive lower burdens than would be implied by discounting their respective expected tax payments at the risk-free rate. Since the total burden in the second period is $(1 + r)B_0$, this means that the sum of the burdens from the individual perspectives will be lower than the present value of the debt repayment—because government policy improves intergenerational risk sharing.

In short, with incomplete risk sharing, we cannot use the valuations of existing generations to discount the flows of future generations. Indeed, we do not even have the valuations of existing generations to rely on for future years that occur after all current generations are deceased.[5]

Our standard approach, then, may overstate the burdens on future generations to the extent that government policy improves intergenerational risk sharing. However, it may be justified with the argument that such benefits of government policy should be considered separately from the first-order redistributions among generations, in the same way that, as discussed in chapter 3, we ignore changes in deadweight loss associated with fiscal policy changes.

In summary, measuring fiscal policy's welfare effects on different generations, as generational accounting seeks to do, requires an evaluation of the risk characteristics of fiscal flows and an appropriate risk adjustment of these flows or, as an approximate substitute, the use of risk-adjusted discount rates. Attempts to sidestep this issue simply by discounting expected flows with a risk-free rate of interest are plagued by the same fundamental problem as deficit accounting—the resulting measures would not be invariant with respect to changes in the superficial labels attached to government transactions. As generational accounting methods to date have not fully identified the appropriate

5. Adding the possibility of incomplete *intra*generational risk sharing would simply extend the complexity one additional step. Even within a generation, the total burden might be lighter than would be implied by discounting that generation's overall payments with a market discount rate.

adjustment for risk, it remains standard practice to estimate generational accounts for a range of discount rates.

References

Auerbach, Alan J., Bruce Baker, Laurence J. Kotlikoff, and Jan Walliser. 1997. Generational accounting in New Zealand: Is there generational balance? *International Tax and Public Finance* 4, no. 2 (May): 201–28.

Auerbach, Alan J., Jagadeesh Gokhale, and Laurence J. Kotlikoff. 1991. Generational accounts: A meaningful alternative to deficit accounting. In *Tax policy and the economy,* vol. 5, ed. D. Bradford, 55–110. Cambridge, Mass.: MIT Press.

Diamond, Peter. 1996. Generational accounts and generational balance: An assessment. *National Tax Journal* 49, no. 4 (December): 597–607.

This Page Intentionally Left Blank