

GENERATIONAL ACCOUNTS AND GENERATIONAL BALANCE: AN ASSESSMENT

PETER DIAMOND *

Abstract - *Generational accounting is presented by Auerbach, Gokhale, and Kotlikoff for normative consideration of government and positive analysis of individuals. Generational balance is presented as a normative criterion and a useful target for political economy. This assessment distinguishes cost and utility bases for such calculations and concludes that the cost basis is interesting and more useful—easier to get straight conceptually and consistent with a residual calculation of the impact on future generations. Generational balance as currently calculated does not seem good by itself for identifying either equitable policies or good ones.*

INTRODUCTION

The development of generational accounting in the joint and separate writings of Auerbach, Gokhale, and Kotlikoff (for example, 1991, 1994, 1995; AGK, hereafter) has three strands. One strand is the role of generational

accounting as part of normative consideration of government activity. A second strand is the use of generational accounting in a positive theory of individual behavior. A third strand is the presentation of generational balance as both a normative criterion and a useful target for political economy. Relative to all of these is also an attack on the conventional deficit (in its standard variations) as a useful tool for economic analysis. In this assessment, I draw on the previous reviews by Baker (1995), Cutler (1993) and Haveman (1994).

THREE PARALLELS

It is useful to start with calculations that bear some partial resemblance to generational accounts and that have some familiarity. The three to be discussed are the projections for Old-Age, Survivors, and Disability Insurance (OASDI) and Medicare, the calculations of the incidence of taxes, and the calculations of federal government impact by state.

Each year projections are prepared for OASDI and Medicare. Since I am more familiar with the former, I will concentrate on that one (Board of Trustees,

*Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02139.

1995). Three projections are prepared, a central “best” projection and two that straddle it to represent some measure of a range of alternative projections. The projection extends 75 years into the future. The focus of the projection is the state of the trust funds. The basis of the projection is existing legislation, ignoring a non-negativity constraint on the trust funds. Being a projection, the calculation describes what would happen to the trust funds if the assumptions behind the projection came to be realized. The central idea is to identify with considerable lead time when the current legislation is likely to be inconsistent, to need changes in benefits and/or revenues in order to have a path on which Social Security can continue. In addition to projecting current legislation, the same methods are used to calculate the implications of alternative legislated changes that might restore actuarial balance. There are a variety of short- and long-run concerns that the projections are meant to flag as reasons for concern about the future of the program.

Since the late 1940s, there have been repeated calculations of how the U.S. tax system affects different groups in the economy (Musgrave et al., 1951; Pechman, 1985). The focus of the analysis is to identify groups by income level and to allocate the tax revenue collected in a year to the different groups, recognizing the incidence of the taxes, not just the legal liability to pay. In addition to the allocation of the entire tax revenue, the methodology is used for differential incidence—the examination of the distributional impact of proposed tax changes.

While the tax analysis considers only tax collection (possibly net of cash transfers), another approach is to consider the impact of the entire government

budget. This is done by identifying groups (for example, by state) and allocating both expenditures and taxes to the different groups (Forsythe and Friar, 1993).

In allocating expenditures, there are two ways that such a calculation could be done, based on costs or utilities. One way would be to take the cost of different expenditures and allocate the cost to different people who are identified as benefiting from the expenditures. This necessarily involves a somewhat arbitrary allocation of common costs. It has the adding-up property that the total allocated is equal to the aggregate of government expenditures. This approach is based on measuring from the government’s perspective who the money is spent on. A different approach would be to calculate a utility-to-individuals for the different expenditures. Thus, the marginal rate of substitution between the government program and income forms the basis of the valuation, rather than a fraction of the cost of providing the program. In this case, one would be trying to value the net utility impact of the combination of taxes paid, transfers received, and benefits received as a result of the expenditures. Note that if there are insurance components to the government-provided transfers, then an *ex ante* utility-based approach to taxes net of cash transfers needs to assess the premium people are willing to pay for the insurance. When everyone is trading the equivalent of all of the government taxes, transfers, and expenditures at the same prices, then everyone will have the same marginal rates of substitution between their net receipts of dollars or benefits from different programs. In the more realistic case that there are not complete and perfect markets, marginal rates of substitution will differ across people.

Depending on how people valued the expenditures, the sum of utilities from the expenditures might be larger or smaller than the aggregate level of expenditures. This difference between the cost and utility approaches in the presence of an adding-up condition is relevant for the ability to allocate a residual aggregate cost to some group rather than estimating directly. That is, with a cost-based approach, having allocated some of the cost to some groups, the cost relevant for the remainder of the population is the difference between the total and what has been allocated. The same residual approach cannot be used with a utility-based approach. In addition, restriction to a single year makes the treatment of investments somewhat awkward, since these are meant, in part, to affect people (somewhat the same people, somewhat different people) in future years.

In all three of these examples, the response of equilibrium in the economy to the programs is not part of the calculation. That is, there is a partial equilibrium nature of the calculation, rather than general equilibrium determination of the full set of changes in the economy.

From the perspective of these examples, generational accounts are prospective and cover many years, like the Social Security projections. Like the other two calculations, generational accounts are concerned with an allocation among groups (identified by date of birth and gender) not in the total budget or public debt. The current version of generational accounts is neither a calculation based solely on taxes (or taxes net of cash transfers) nor a calculation allocating the full government budget. Rather, it is the somewhat anomalous combination of taxes, cash transfers, Medicare, and Medicaid (and, in some other

countries, educational expenditures), an anomaly related to the definitions in the National Income and Product Accounts (NIPA). The current calculation is a mix of cost-based calculations and utility-based calculations, which again is anomalous. In terms of the grouping of the population, all the not-yet-born are combined using a residual calculation, not separate generational calculations. The lack of a consistent cost-based calculation seems to make the use of a residual calculation for future generations inappropriate.

GENERATIONAL ACCOUNTING— RETROSPECTIVE

It seems to me that we are interested in comparing the impacts of alternative government policies on people born at different times. Thus, we want some generational accounting, although we also continue to want calculations based on other groupings. The issue in this section is what generational accounting can hope to measure and how it has been and could be done.

To delay some issues about projection, I begin by discussing historical accounting. Assume that one had a full history of government actions. Year-by-year and person-by-person, assume that we had a complete description of the impact of government actions. (For some people in some settings, even the year may be too long an accounting period—for example, the utility of poverty alleviation measures will depend on the frequency and timing of payments.) For any individual, we would see the set of tax payments to the government. Given an incidence theory, one would also have year-by-year estimates of taxes “really” paid. Similarly, one would have a series of year-by-year estimates of cash transfers received. One also wants to value the benefits received other than in

cash. (In a world with military drafts, there are also costs-in-kind as well as benefits-in-kind.) This would include goods received free (highway use, for example) and goods sold at less than what they would have cost without the government action (access to national parks, for example). One might want to recognize individual uncertainty and calculate as of some date the distribution of transfers that might have been received at some later date, given the rules in force when that later date happens—as with annuitized retirement benefits. Indeed, evaluation of insurance programs on an individual *ex post* basis would be peculiar.

It is worth noting that similar to the contrast between evaluating goods and services on a cost or utility basis is the contrast between evaluating cash payments that contain an insurance component on a cost or utility basis. Since people pay markups when they purchase insurance privately, any utility-based evaluation of the government provision of insurance should involve such a markup. For example, the real annuities provided by Social Security should be worth considerably more (relative to expected cost) than the nominal annuities in the market, which themselves show substantial markups. Thus, a cost-based calculation of Social Security uses the amounts actually paid (possibly marked up by the small administrative costs). With a utility-based calculation, the benefits received would be considerably marked up, multiplied by a factor possibly reaching 1.5. This is a nontrivial issue. By making no markup, current generational accounting is using a cost or, more accurately, a cash-received basis.

This issue can also be seen by considering a government program that

precisely displaced identical private insurance, having no real effects whatsoever. Assume the government program is identical to the displaced private insurance, including the costs of running the program. Then, examination of the cash flows would find that individuals are losing from the program to the extent of the administrative costs, even though real allocations are unaffected. In practice, Social Security provides insurance that is different from that in the private market and has administrative costs that are lower than those of private insurance markets.

The next step, assuming one wanted to take it, is to determine how to convert this array of annual estimates into a scalar. While this is simple in a world with complete markets that are available to all at the same prices, the issue is a mess when different people face different interest rates, and incompleteness of markets implies that people are making personal expenditure decisions under uncertainty about future transaction opportunities and future resources and needs.

While tempting, it is inadequate to collapse this issue into a choice of interest rate. The inadequacy parallels the same issue in benefit-cost analysis, where the determination of shadow prices for government expenditure decisions is not adequately discussed as a debate over the right discount rate. (For a discussion of tax incidence in the presence of liquidity constraints, see Hubbard and Judd, 1986.) It is the interest rates on government debt that tie together the histories of tax revenue, expenditures, and debt in a single budget constraint. Assuming the tax receipts from the taxation of interest income are separately recognized as part of taxation, it is gross-of-tax interest rates paid that become the right

measure for evaluating the net cost-to-government of dealing with people in different years. On the other hand, the interest rates relevant for evaluating the impact on individuals are generally different interest rates (apart from the people who are holding such debt throughout their lives, assuming that the interest was not taxed and that these rates represent marginal decisions). And the choice of interest rate is different for different portions of the population. This is particularly an issue with Social Security. If you think that a major justification of Social Security is paternalism—to force people to save more than they would if left to themselves—then it would be potentially misleading to value Social Security using the interest rates inherent in the preferences of the people subject to its mandates. On the other hand, this can be viewed as a cost of paternalism where a mandate is deemed inappropriate. However, one cannot “get off” that easily if the interest rates relevant for combining cash flows in two different years are different when perceived at alternative earlier dates. That is, at age 21, I might make a decision that affects relative incomes at ages 50 and 60 in a way that reveals an implicit interest rate. At age 40, I might make a decision that shows a different implicit interest rate. There is no obvious answer to what the historian should do in making this calculation. This problem arises from incomplete markets. The same problem arises if preferences for the future change over time. In this case, it is hard, and possibly impossible, to develop the normative basis for a single lifetime utility-based calculation.

For a single interest rate, the same for all members of a generation, AGK choose a rate meant to reflect what they estimate to be the return equal to the rate required to hold an asset with

risk characteristics approximating those of government taxes and spending. This is a focus on lifetime utilities, not the cost of government. This is not the interest rate which satisfies their present discounted value government budget constraint. Since many people hold little or no assets, this is an approach that makes some sense for the well-off, not for the population in general. Nor does it correct for taxes on interest earnings which would be appropriate for a utility-based approach. Perhaps, consistent with this logic, different interest rates should be used for different government programs, with much higher interest rates used in part. But then, we know that it is not generally true that the correct adjustment for risk is simply a different interest rate, that then gets compounded repeatedly into the future. They do report calculations with different interest rates, so one can see the effect of different choices. Using an interest rate above the government rate and calculating Social Security benefits on a cash-received basis is combining a utility-based calculation with a cost-based calculation.

Of the assumptions used by AGK, I have a few comments on the treatment of capital income taxation. They capitalize the inframarginal fraction of part of the income, and all of the corporation and estate taxes, but do not capitalize any of the taxation of residential capital. While some benefits of local government also get capitalized, suggesting some netting out, a more detailed treatment seems warranted. In determining the fraction of the personal income tax to multiply by the fraction inframarginal in order to know how much to capitalize, they take the fraction of capital income in Net National Product (NNP). This seems odd since capital income is so heavily received by high-income people, while much of labor income goes to people not subject to income taxation.

GENERATIONAL ACCOUNTING— PROJECTION

Contrasting the future with the past, there are two further issues. One is to project the economic environment for the government. For example, how much will government-provided medical services cost and how much will defense cost. Second is to describe what should be used as the course of government action. That is, what projections should be made about future legislation, either to extend the current way of doing things into the future or to recognize that some ways of doing things are likely to be changed. This is a major issue since many projections of current practices lead to politically unlikely or even inconsistent outcomes. For example, both the Medicare and OASDI trust funds are projected to go negative if existing legislation on taxes and benefits is projected into the future. What projection should be used in a setting like this? The answer depends on the purpose for the accounts. If one wants to show the inconsistency of current Social Security legislation, one wants to project current legislation. If one wants to show the detailed generational impact of that inconsistency, one would need to select an alternative source of revenue to maintain the promised benefits or an adjustment of the benefits. (Conceptually simpler is the comparison on a generational basis of alternative methods for restoring actuarial balance.) Alternatively, for positive analysis of short-run private behavior, the projection should be based on the beliefs of the private agents. In considering how generational accounts for the entire budget should reflect Social Security, there is a similar dichotomy. If the purpose is to identify the concerns with general revenue and general revenue financed programs, then one might

assume that Social Security will receive new legislation that restores actuarial balance. One way to allow for the need for such an adjustment would be to leave Social Security out of the generational accounts. This is different from the situation in which one considers the impact of all government programs. A similar issue arises with the integration of federal, state, and local accounts. This involves an implicit political assumption, when one is using the accounts to identify future federal revenue issues.

For projection purposes looking for inconsistency, describing a scenario that might happen, the evolving budget constraint that will affect future taxes depends on the realized real interest rates. *Ex ante* combinations of different stochastic future cash values are not part of such a projection of a scenario. That is, for some purposes, one might want a projection of what will happen if a particular set of economic and government behavior assumptions is met. For other purposes, one might want an expected utility calculation of how the entire future is perceived. The latter involves, explicitly or implicitly, a behavior of government as a function of how the economy evolves. For example, one would evaluate Social Security differently depending on whether one viewed future legislation as zero-sum political risk or as a device for intergenerational risk spreading in the presence of incompleteness in the “contract” covering the relationship between Social Security and the evolution of the economy.

After 2004, the 1994 Office of Management and Budget (OMB) projection assumes that expenditures per person in three broad age groups will increase with productivity. Thus, the projections assume that changes in the relative sizes

of the groups result in no changes in per capita expenditures. This might overestimate the effects of population change on spending since expenditures per capita such as education might adjust to the size of the receiving population, with a larger group receiving less per person. This assumption is important in the comparison of how generational accounts vary with the rate of productivity increase. However the projection is done, the variation of outcomes with the rate of productivity growth is sensitive to the relationship between the projection of expenditures and the assumed rate of productivity growth.

POSITIVE GENERATIONAL ACCOUNTING

Starting with a life-cycle model of individual savings, the authors view the calculation of changes in a lifetime budget constraint as a good basis for predicting the private market response to new legislation. Liquidity constraints make this highly questionable as an approach to short-run behavior. For example, consider the government introduction of mandatory loans to the government (that are not tradable and not usable as collateral), paying the market interest rate. Since liquidity constrained individuals cannot offset this reduction in their access to purchasing power, this has an impact on short-run demand. Having had an effect on short-run behavior, this policy then has an effect on the later long-run behavior of the same people; as a consequence of any short-run impacts, people reach older ages with different wealth than they would otherwise have had. This can affect their choice of planned bequests and so their lifetime marginal propensity to consume. It is clear that the lifetime marginal propensity to consume scalar used in Gokhale, Kotlikoff, and Sabelhaus (1994) is not a sufficient statistic for many government

activities, including the annuitization provided by Social Security and Medicare.

Another issue arises with the role of actual legislation. The AGK calculation is based on a projection of government policies meant to reflect continuation of current policies. Insofar as the behavior of individuals is based on projections of government policies, it will be based on the projections by the individuals of what they think might actually happen. These two approaches to projection differ in a central way. Insofar as generational accounting is meant to highlight the unsustainability of continuing current policies, it is not a good instrument for estimating individual responses. Individuals will focus on what they think will happen, not what would happen if current unsustainable policies were unchanged. For example, when the government projects the inconsistency of current Medicare and OASDI legislation, we would expect individuals to recognize the possibility of change in making their savings decisions. Thus, both the 1977 and 1983 Social Security changes had massive decreases in projected future benefits. However, current consumption behavior did not reflect a massive drop in lifetime income when the legislation passed, since the nonviability of existing legislation was clear.

GENERATIONAL ACCOUNTING—TELESCOPED

AGK allocate net taxes (or taxes and expenditures) to current generations. While one could use the same methods to make allocations to future generations one at a time, AGK do not do this. Rather, they use the government budget constraint and the assumption that each future generation bears the same net cost relative to income in order to

calculate a single net (per generation) cost for all future generations. This is different from projecting the current policies indefinitely into the future. Instead, it is a residual calculation of a net amount available across the birth dates of these different generations.

Among the questions raised by this procedure is the role of the interest rate in collapsing these future net costs into the present. Insofar as one is taking a cost basis, then the government interest rate is the appropriate one to use and there seems little ambiguity (for a projection, as opposed to recognizing explicitly the uncertainty about the future). However, if one takes a utility-based approach, then one uses marginal rates of substitution in order to infer the appropriate interest rate. But for any given generation, these are only inferable over the period that the generation is alive; it is not a basis for selecting interest rates for the time before their birth. In addition, the government budget constraint holds for government cash flows, not for the utilities associated with those cash flows. Thus, I think that one cannot telescope as AGK do with a utility basis, but only with a cost basis. In the vocabulary used above, there is no adding-up constraint on a utility basis that makes the sum of utilities equal to any constant. One cannot then infer the net cost to future generations as a residual.

This telescoping is part of the procedure AGK follow on the way to examining generational balance. One has no need for telescoping for some interesting calculations, such as what size (permanent) increase in the income tax would yield government present-discounted-value budget balance given a projection of current expenditure patterns. This is analogous to the calculation of the size

of the Social Security actuarial imbalance relative to taxable earnings. Yet such a calculation based on the income tax ignores possible changes in FICA taxes that might happen to restore actuarial balance in Social Security.

GENERATIONAL BALANCE—NORMATIVE

In addition to generational accounting, calculating the lifetime economic impact of government actions, AGK consider generational balance. Generational balance is defined as equality between the growth-adjusted lifetime accounts of current newborns and future generations. For current newborns, the calculation is done using the projections of government policies; for future generations, the calculation is done as a residual. Calculating the ratio of net taxes to projected income, constant for all future generations, that will balance the government's infinite horizon budget. As argued above, this calculation only represents a constraint on the government if the calculations are done on a cost basis, including the interest rate faced by the government. Thus, a utility-based calculation, which might be seen as the most interesting from a social welfare function perspective, is not consistent with this residual calculation, which has a cost basis.

There are three questions to ask about generational balance. First, assuming the calculations are done fully, is balance in a utility-based approach a social welfare optimum. Second, if the calculation is done on a cost basis, does the level of normative value extend to this calculation. (I do not ask about the mixed calculation with only some expenditures included.) Third, does calculating what it would take to restore balance have a valuable political role, even if there is not a social welfare basis for choosing this outcome precisely; that

is, the distinction between generational balance as a normative criterion and as a political tool.

It is clear that an optimum overlapping-generations (OLG) growth calculation will not generally have generational balance on either a cost or utility basis. Redistributing across generations will be part of such an optimization with probability one. Second, one would want to look at a full calculation of the impact of the government, not a partial one. That is, one would not draw overall normative conclusions from consideration of part of government activity. To indicate an example of the problems run into with a mixed calculation, the Appendix contains a model where the economy is in a steady state, but a calculation would show imbalance.

GENERATIONAL BALANCE—POLITICAL

I conclude that the role of generational balance is in the political economy realm, not that of normative economics. From this perspective, the question is the utility of going beyond the use of generational accounting to assess the differential incidence of alternative government policies and calling for generational balance. Two questions arise with any such political calculation—the extent to which a call for generational balance would help to pressure government to have a wider perspective of the effect of policies than would otherwise occur and the extent to which unfortunate policies would be pursued because they make a summary statistic look better, without having the same effect on the underlying reality. Focus on the conventional deficit encourages government actions that look good in this light, even if they are not necessarily helpful with capital accumulation, which is a major part of the concern with the

deficit. One example is the sale of assets.

Exclusive attention to generational balance would also have its limitations. One example would be the impact on balance of policies far into the future which might be legislated relatively painlessly since they would have no visible effects in the near term. This reflects the tension between short horizon measures which encourage substitution across the time limit in the measure and long horizon measures which allow announcement of deferred policies, which might be better implemented sooner and also might not happen.

In addition, since generational balance is a comparison of newborns with the future, there are policies that are clear changes in generational distribution that would not show up as a change in generational balance. For example, a program of increased taxes used for education would affect newborns and future generations the same and so would not contribute to generational balance even though current older generations are paying higher taxes for the benefit of future generations.

To see an example of a problem with generational balance considering only taxes net of cash transfers, consider the role of assets that provide services that are given away, such as the interstate highway system. AGK value government wealth by capitalizing fee income received by the government at an interest rate. As they point out, a complete adjustment in this framework would include the value of physical assets deducted from the public debt and the future services provided by these assets added to the government expenditures to be financed. Since the framework does not value the services,

but only recognizes the cost of providing them, this is an inadequate picture of intergenerational fiscal relations and would have a perverse political incentive from a drive for generational balance. The perversity comes from the equal contribution to generational balance from cutting expenditures on current public consumption and cutting expenditures on current public investments that do not generate revenue.

Including some but not all government programs raises problems wherever the line is drawn between what are identified as transfers (and included) and what are not. Consider a program to give vouchers for school lunches as opposed to giving grants to schools to provide lunches. These are economically very similar in their generational distributional consequences. Yet would one count the former in taxes net of transfers but not the latter?

Conclusions

I conclude that generational accounting is interesting. It is interesting whether one looks at future taxes or future taxes net of future expenditures, and it is interesting whether one looks at future generations and the trajectory of the public debt separately or telescopes them through the government budget constraint. At a minimum, I think it would be beneficial to have official projections well out into the future, paralleling those of the Social Security Trustees. This might lessen the political tendency to overfocus on a particular period, such as the year 2002, leading to misleading statements about some policies. Also, I think it would be beneficial to devote more research resources to examining how to do generational accounting better and to report alternative calculations—AGK

have been laboring without as much feedback from the rest of the profession as they deserve. I think that current focus on the cost basis for projection is likely to be more useful than current focus on the utility basis. This approach is easier to get straight conceptually and is consistent with the use of a residual to calculate the impact on future generations. I do not think that generational balance calculated as it is currently is a good basis by itself for identifying either equitable policies or good ones. Of course, neither is an exclusive focus on conventional budget balance a good basis for identifying either equitable policies or good ones.

ENDNOTE

I am grateful to Alan Auerbach and Larry Kotlikoff for helpful comments.

REFERENCES

- Auerbach, Alan J., Jagadeesh Gokhale, and Laurence J. Kotlikoff.** "Generational Accounts: A Meaningful Alternative to Deficit Accounting." *Tax Policy and the Economy* 5 (1991): 55–110.
- Auerbach, Alan J., Jagadeesh Gokhale, and Laurence J. Kotlikoff.** "Generational Accounting: A Meaningful Way to Evaluate Fiscal Policy." *Journal of Economic Perspectives* 8 No. 1 (Winter, 1994): 73–94.
- Auerbach, Alan J., Jagadeesh Gokhale, and Laurence J. Kotlikoff.** "Restoring Generational Balance in U.S. Fiscal Policy: What Will It Take?" *Federal Reserve Bank of Cleveland Economic Review* 31 No. 1 (1st Quarter, 1995): 2–12.
- Baker, Dean.** *Robbing the Cradle?* Washington, D.C.: Economic Policy Institute, 1995.
- Board of Trustees, Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds.** *1995 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds.* Washington, D.C.: Government Printing Office, 1995.
- Cutler, David M.** "Review of *Generational Accounting: Knowing Who Pays, and When, for What We Spend* by Laurence J. Kotlikoff." *National Tax Journal* 46 No. 1 (March, 1993): 61–7.

Forsythe, Dall, and Monica Friar. "The Federal Budget and the States: The Flow of Funds Between the Federal Government and the States: A Ten-Year Review." Taubman Center for State and Local Government, John F. Kennedy School of Government, Harvard University Working Paper. Cambridge, MA: Harvard University, 1993.

Gokhale, Jagadeesh, Laurence J. Kotlikoff, and John Sabelhaus. "Understanding the Postwar Decline in United States Saving: A Cohort Analysis." NBER Working Paper No. 5571. Cambridge, MA: National Bureau of Economic Research, 1996.

Haveman, Robert. "Should Generational Accounts Replace Public Budgets and Deficits?" *Journal of Economic Perspectives* 8 No. 1 (Winter, 1994): 95–111.

Hubbard, R. Glenn, and Kenneth L. Judd. "Liquidity Constraints, Fiscal Policy, and Consumption." *Brookings Papers on Economic Activity* 1 (1986): 1–61.

Musgrave, Richard, J. J. Carroll, L. D. Cook, and L. Frane. "Distribution of Tax Payments by Income Groups: A Case Study for 1948." *National Tax Journal* 4 No. 1 (March, 1951): 1–53.

Pechman, Joseph. *Who Paid the Taxes, 1966–85.* Washington, D.C.: The Brookings Institution, 1985.

U.S. Office of Management and Budget. "Generational Accounting." In *Budget of the United States Government: Analytical Perspectives, Fiscal Year 1995*, Chapter 3. Washington, D.C.: Government Printing Office, 1994.

APPENDIX: A MODEL

To see the consequences of selecting a higher rate for discounting the future than is paid on government debt, consider a simple model where each generation (all of the same size) lasts one period, so there are no complications from overlapping generations or from differently timed payments for a generation. Assume that in the past the government has borrowed B in order to finance a durable asset that gives each generation a flow of services or to finance a war. Assume that each generation pays taxes, T , equal to rB , to cover the interest cost on the debt. It would appear that each generation was in exactly the same position and there would be intergenerational budget balance. However, if a higher interest rate than the government rate is used in the calculation, there is not balance. The imbalance comes from the fact that taxes are calculated for the current generation, and these are determined by the government interest rate. But a residual calculation using a different rate would result in imbalance.

The current generation pays taxes equal to rB . For consistency, we assume that these are paid at the end of the period and need to be discounted back to today to be comparable to government wealth, which is $-B$, equal to government debt. Thus, its generational burden is $rB/(1+r)$, where r' is the discount rate used in the calculation of generational balance.

Next, we calculate the burden on future generations, denoted N . AGK state that, discounted back to the present, the burden on each future generation is the same, apart from growth, which is assumed to be zero here. This is not precisely what they mean—for the infinite sum in equation 10 in their 1991 paper would imply a zero value of N . They must mean that each generation has the same tax discounted back to its birth, adjusted for growth. In this case, we sum $N/(1+r') \exp(t)$ from $t = 1$ to infinity. This is equal to N/r' .

Thus, the formula for the burden on future generations is

1

$$N/r' = B - PDV(T) = B - rB/(1+r).$$

Thus, the generational balance expression is

2

$$N - rB/(1+r) = B(r' - r).$$

Thus, there is not generational balance even though we are in a steady state with every generation paying the same tax. Clearly, this is an unsatisfactory solution. AGK might (and did) argue that, with certainty in this model, r' should be set equal to r . However, if we introduce the agency costs associated with private loans that are a major contributor to the spread of interest rates in the economy, then we will have this problem back without any aggregate uncertainty. (Formal modeling would need extended lives in order to have loans.) The same point can be made in another way. With a debt of B and an interest rate of r , the government debt can be financed by annual payments equal to rB . If these are discounted back to the present at a different interest rate, r' , equal to the rate used for discounting the taxes of future generations, then the present discounted value of government interest payments is rB/r' , not B . The thrust of AGK's argument appears to be that, since the future is uncertain, the current generation should pay more in taxes than the expected utility of taxes on future generations. This seems unsatisfactory.