

NBER WORKING PAPER SERIES

DISTRIBUTIONAL EFFECTS
ON A LIFETIME BASIS

Don Fullerton
Diane Lim Rogers

Working Paper No. 4862

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1994

We are grateful for suggestions from David Bradford, Gib Metcalf, Joel Slemrod, and Timothy Smeeding. We also received help from many individuals listed in the forward to our 1993 book, as this paper uses the model and data described there. Financial assistance was provided by the American Enterprise Institute and the National Science Foundation grant SES-9122785. This paper is part of NBER's research program in Public Economics. Any opinions expressed are those of the authors and not those of the Congressional Budget Office or the National Bureau of Economic Research.

© 1994 by Don Fullerton and Diane Lim Rogers. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

NBER Working Paper #4862
September 1994

DISTRIBUTIONAL EFFECTS
ON A LIFETIME BASIS

ABSTRACT

All government agencies charged with the responsibility of estimating distributional effects use annual income to classify households and one year's tax to characterize tax burdens. In this paper, we describe an alternative procedure to estimate lifetime tax burdens as proportions of lifetime income. To illustrate this model, we calculate lifetime effects of a uniform consumption tax and a wage tax. This kind of analysis can supplement existing annual analyses, since policymakers might want to insure both that current taxes reflect current ability to pay and that lifetime taxes reflect lifetime ability to pay.

Don Fullerton
Department of Economics
University of Texas at Austin
Austin, TX 78712-1173
and NBER

Diane Lim Rogers
Congressional Budget Office
Second and D Street, SW
Washington, DC 20515

Distributional Effects on a Lifetime Basis

I. INTRODUCTION

Recent analyses of tax reforms have placed growing emphasis on determining the distributional effects of policies. Other papers in this volume address various issues concerning the distributional tables used by all government agencies charged with the responsibility of measuring distributional effects of taxes. As discussed in these papers, agencies disagree about the unit of analysis, the measurement of income, tax incidence assumptions, and treatment of behavioral effects. In measuring both burdens and abilities to pay, however, all of the agencies adopt a relatively short-run time horizon despite a recent emphasis in the academic literature on lifetime measures. The Congressional Budget Office (CBO) has sometimes conducted analyses using consumption as a proxy for long-run income, and the Joint Committee on Taxation (JCT) has recently adopted a 5-year horizon for their calculations, but none of the agencies have embraced a full lifetime horizon. Categorization of households strictly according to annual income is still the norm. This apparent divergence between policy application and academic emphasis is attributable to concerns both philosophical and practical in nature. The various arguments against using the lifetime framework, and our responses to these criticisms, are discussed later in this paper.

We argue here that lifetime models of tax incidence can play a valuable role in the analyses of real-world tax reforms. The concept of the lifetime incidence of a tax simply suggests an additional examination, namely, of lifetime tax burdens relative to lifetime incomes. Our model generates distributional tables very similar to those which are already widely used in government circles. Because lifetime income may not reflect current ability to pay taxes, however, the lifetime framework can a supplement rather than replace the traditional annual approach. In other words, policymakers might worry that current taxes reflect current ability to pay, and that lifetime taxes reflect lifetime ability to pay.

The differences between annual and lifetime incidence have considerable bearing on certain policy issues. A good example is the ongoing debate about the relative merits of consumption-based taxes versus income-based taxes. Proponents stress that a consumption tax would increase economic efficiency by removing the double taxation of interest income. Dynamic life-cycle models can calculate the welfare gains from removing these intertemporal distortions, so the lifetime perspective is likely to favor consumption taxes from an efficiency standpoint. Opponents stress that consumption taxes are regressive, placing relatively greater effective tax rates on households with low income. Yet this latter argument is based on an annual perspective. A consumption tax must be regressive with respect to annual income, since consumption makes up a high fraction of earnings for low-income groups. From a lifetime perspective, however, Poterba (1989) and others point out that if the present value of lifetime consumption equals the present value of lifetime income, then the lifetime incidence of a uniform-rate consumption tax is necessarily proportional. Thus consumption-based taxes are likely to look much more desirable when using a lifetime model than when using an annual model, from both efficiency and equity standpoints.

A wage tax has steady-state effects much like a consumption tax, both in terms of efficiency and in terms of lifetime incidence. It differs in transition, however, since retired generations would prefer the switch to a wage tax rather than the switch to a consumption tax.

In this paper we use a consumption tax and a wage tax as examples for measuring lifetime distributional effects, and we generate results that shed new light on the tradeoff between intertemporal efficiency and distributional equity. We use a sophisticated general equilibrium model developed in Fullerton and Rogers (1993) to calculate the economic incidence, that is, the distribution of real burdens from a consumption tax. Most importantly, we consider the regressivity of a consumption tax when individuals are classified not by their annual income but instead by lifetime income. Distinguishing among different lifetime-income categories, we are

able to characterize the pattern of lifetime tax burdens relative to lifetime incomes, using utility-based "equivalent-variation" measures.

Within the model we have developed, the lifetime incidence of a uniform consumption tax is a bit more interesting than the purely proportional outcome suggested by Poterba. In our model, for three reasons, a consumption tax is typically not proportional on a lifetime basis. First, some groups receive more bequests which enable them to spend more than they earn. Second, individuals are classified by the present value of potential earnings, as a measure of well-being, so some groups may choose to work more and consume more taxed commodities instead of untaxed leisure. Third, the distributional effects of a consumption tax depend on which current tax it replaces.

To concentrate on these issues regarding inherent regressivity, we consider a flat tax on all consumption at the same rate. This uniform consumption tax can be interpreted as a proportional tax on individual consumption, as a national sales tax on all goods, or as a value-added tax (VAT) at a single rate.¹ We thus ignore many differences in the administration of the tax, especially whether it is collected from individuals or from firms. And we ignore any progressivity that might be added into the system, either through graduated rates on an individual consumption tax base or through exemptions of necessities from a value added tax base.

The next section describes current approaches that categorize households either by annual

¹An individual consumption tax could operate like the current individual tax, with the inclusion of all sources of income and the deduction of all forms of savings. Since all income must either be consumed or saved, income minus savings must equal consumption. This tax base can be applied to a progressive rate structure. In contrast, a VAT would apply to value-added at each stage of production, measured by sales revenue minus the cost of all intermediate inputs. If full deductions are allowed for the cost of capital goods, then the VAT is a consumption-based tax. On each final product, the total revenue collected from a VAT is equal to the sum of amounts paid at each stage. Since the value of the final product is just the sum of value-added at each stage, this consumption-based tax is equivalent to a retail sales tax. A specific VAT proposal may exempt certain goods and place different rates of tax on different industries. We do not consider these many variants.

income or by age, and how we combine those two approaches. We outline our efforts in five steps, and we summarize results. Section III then addresses a fundamental obstacle to the use of lifetime models in distributional analysis, namely, the lack of data on lifetime incomes. We explain how we tackle this problem by econometric estimation of lifetime incomes using panel data. Section IV describes the general equilibrium nature of our model, which allows us to account for behavioral responses and price changes in calculating the distributional effects of tax policies. Greater detail on our model can be found in Fullerton and Rogers (1993). In Section V we simulate the replacement of each major U.S. tax with a uniform consumption tax, and with a wage tax, and we determine the lifetime regressivity or progressivity of the various existing taxes relative to these two alternatives. Distributional tables show lifetime equivalent variations (tax burdens) relative to lifetime incomes (lifetime abilities to pay) under two parameterizations. Section VI discusses some of the arguments against incorporation of lifetime analysis into the policy arena. Section VII concludes.

II. HOW WE COMBINE EXISTING APPROACHES

Existing studies of tax incidence fall into two general categories. The most common approach in policy analysis begins by dividing all households into groups based on some measure of their current annual economic income.² Such studies then collect data on the wages and salaries of households in each group, capital income such as interest and dividends, and expenditures on each commodity. Shifting of tax burdens onto each group can then be measured from estimated changes in wages, interest rates, and commodity prices.

A problem with this approach, however, is that the lowest-annual-income group is a mixed bag. It includes some young workers just starting a career who are expected to earn more

²This approach is exemplified by Pechman and Okner (1974) and Ballard et al (1985).

later, some retirees who had earned more earlier, some people with volatile income who just had a bad year, and, finally, the perennially poor. Yet the concern of policy must differ for these four types of individuals: some are really poor, some are comfortable, and some are actually very rich. Even if careful study shows that a particular tax change redistributes from a high-income group to a low-income group, little is known about what happens to the welfare of individuals who move among the annual income classes during their lives. Indeed, most of us move up the income scale during working years and then back to low income upon retirement.

The second approach proceeds by dividing all households into groups on the basis of their age. Such studies are typically based on the "life-cycle" model, that is, a specification of individuals' entire lifetime plans for saving and consuming during working years followed by dissaving and consuming during retirement years.³ Such a model is used to calculate equilibrium prices over time and to report present value changes in lifetime tax incidence. These studies focus on how taxes affect savings incentives, capital formation, and future productivity. They can also simulate a tax change, estimate effects on wage rates and interest rates, and measure redistributions between young and old.

A simpler, partial-equilibrium version of this second approach is used in the "generational accounting" framework discussed by Laurence Kotlikoff in this volume. This framework focuses on what the government's intertemporal budget constraint implies about the intergenerational incidence of government policies. Effective tax rates can be computed for many different cohorts of households. For simplicity these measures assume constant prices and unchanged behavior, but the consequence is that generational accounts for alternative policies do not reflect predicted outcomes or changes in economic welfare. Generational accounting seems most valuable in addressing questions about current policies, such as the extent to which each

³Examples are Summers (1981) and Auerbach, Kotlikoff and Skinner (1983).

generation pays for its own programs over its own lifetime -- an assessment that annual deficit accounting clearly fails to provide.

A general problem with this second approach is that it considers only one kind of individual in each age group. Different generations are alive at one time, and the model can calculate effects on these different age groups, but individuals are not distinguished by level of well-being. This kind of model misses the fundamental distinction between rich and poor, a distinction that plays prominently in any policy debate about the distributional effects of taxes.

To illustrate the distinction between the two approaches, suppose that the economy included only the two types of individuals depicted in Figure 1. One type has relatively poor lifetime prospects, advancing with age through points A, B, C, and D. The other type has relatively rich prospects, and advances with age through points E, F, G, and H. The typical annual incidence study would take individuals at point G as the highest-income group, lump together individuals at points F and C for the second group, lump together those at points E, B, and H for a third group, and add those at points A and D for the poorest group. They might then find how taxes redistribute among these groups, but the results would convey nothing about what happens to either of the two individuals of concern.

In contrast, the typical life-cycle or generational accounting study would lump together individuals at points A and E as one youngest group, those at B and F as another group, C and G as the next group, and D with H as the oldest group. The model could then calculate redistributions between the old, the young, and later generations, but not between rich and poor. Neither of these approaches captures the fundamental distinction between the two types of individuals in this economy.⁴

⁴Just a few studies have considered the incidence of taxes across different lifetime income categories. Using Canadian data, Davies et al. (1984) construct lifetime histories of earnings, transfers, inheritances, savings, consumption and bequests. They then measure lifetime income, classify households, and add up the burdens under alternative incidence assumptions. Thus, they

Our new model combines these two approaches. We distinguish between rich and poor, but we classify individuals on the basis of lifetime income. We use the model here to evaluate the lifetime incidence of a consumption tax and a wage tax. In order to avoid the debate about deficit reduction, however, we simulate the effects of using a consumption tax or wage tax to replace each existing U.S. tax, one at a time. This "differential incidence" depends on features of the tax being replaced, so we must specify details of personal income taxes, corporate taxes, and all other U.S. tax instruments.

This effort can be described in five major steps. First, we start with data from the Panel Study of Income Dynamics, including thousands of observations of individuals over an 18-year period. With all individuals together, we estimate the wage rate as a function of age and other demographic variables. We then construct a lifetime wage profile for each individual, using actual wage rates for available years and predicted wage rates in other years. We use this wage profile to calculate the individual's present value of potential lifetime earnings. We rank individuals by lifetime income and classify them into twelve lifetime income groups.

Second, for each lifetime income group, we re-estimate the wage profile as a function of age. We also estimate age-profiles for the personal income tax and for government transfers such as social security. The estimated wage profile rises and then falls over time for all groups, but the steepness and the timing of the peaks differ. Groups that earn relatively more of their income earlier in life must save more for retirement. These groups are then likely to benefit more from a consumption tax used to replace any tax on capital income. Also, we add to each

use the same basic approach as Pechman and Okner but extend it to a lifetime context. Poterba (1989) focuses on sales and excise tax burdens in the United States. Appealing to the permanent income hypothesis of Friedman (1957), he uses current consumption as a proxy for lifetime income in order to classify households. He agrees with Davies et al. that sales and excise taxes are less regressive in the lifetime context. Lyon and Schwab (1990) use data from the Panel Study of Income Dynamics in a model of life-cycle behavior to find that cigarette taxes are just as regressive, and alcohol taxes are slightly less regressive, when measured with respect to lifetime income rather than annual income.

group an estimate of inheritances. These are highly concentrated at the top of the income distribution, and they also affect the incidence of replacing taxes on capital income.

Third, we use data from the Consumer Expenditure Survey to estimate how people allocate their consumption among specific commodities. These choices depend on age and on income. If a uniform consumption tax replaces one that affects relative product prices, it may then burden some groups more than others.

Fourth, we build a general equilibrium simulation model that encompasses all major U.S. taxes, many industries, both corporate and noncorporate sectors within each industry, and consumers identified by both age and lifetime income. It is not a model of annual decision-making, but a life-cycle model in which each individual faces a particular inheritance, a set of tax rules, a wage profile, and a transfer profile. Each then plans an entire lifetime of labor supply, savings, demand for each consumption commodity, and bequest. We also specify producer behavior, which determines each industry's use of labor, capital, and intermediate inputs. We can then simulate the effects of a tax change on each economic decision through time. We calculate new labor supplies, savings, capital stocks, outputs, and prices. With effects on all ages in all years, we can also calculate the change in economic welfare for groups ranging from those with low lifetime income to those with high lifetime income.

Fifth, we evaluate the effects of using a consumption tax or a wage tax to replace each U.S. tax. In our model these two replacement taxes are not equivalent. During the transition, retired individuals prefer the wage tax. Even in the steady state, because of inheritances, the present value of consumption for any individual exceeds the present value of wage income.⁵ Thus the consumption tax base is larger than the wage tax base, allowing a lower tax rate and

⁵Individuals are born with an endowment of capital and must die with the same amount of capital, augmented by the rate of economic growth. Because the specified growth rate is lower than the specified interest rate, consumption possibilities are augmented by these bequests.

a greater efficiency gain. Also, because these inheritances are concentrated in the higher lifetime income categories, the consumption tax is more burdensome than the wage tax on the lifetime rich.

Using either consumption or wage taxes, we find that replacement of the personal income tax is regressive, while the replacement of current sales and excise taxes would be progressive. The reason is that current sales and excise taxes place higher rates on goods like alcohol and tobacco that constitute a higher fraction of low-income groups' spending. Replacement of the payroll tax is also progressive, because the current tax is regressive.

Replacement of the property tax raises the net rate of return to capital in our model, so it helps high-income groups on the sources side. It also reduces the cost of housing, however, so it benefits low-income households on the uses side. Thus the overall pattern of gains is U-shaped across our lifetime income groups.

Our results for replacement of the corporation income tax are a bit surprising. In the standard analysis, replacement of the corporation income tax would raise the net rate of return to all owners of capital. Since these capital owners are in high-annual-income brackets, the change would be regressive. In our base year of 1984, however, the corporation income tax collects very little revenue.⁶ Any tax that might have been collected on the return to equity is largely offset by interest deductions, investment tax credits, and accelerated depreciation allowances. The removal of the corporate tax system of credits and deductions hardly affects the overall rate of return on the sources side. It does raise costs for industries that received more than the average amount of investment tax credits and accelerated depreciation allowances, however, and it therefore reduces the relative cost of other outputs such as tobacco and gasoline.

⁶This revenue is calculated from observed capital stocks, an assumed 4 percent net rate of return, and effective tax rates that reflect the statutory tax rules for different assets under 1984 law. It thus reflects a long-run equilibrium, not short-run profit fluctuations.

These goods constitute a high fraction of low-income budgets, on the uses side, so the change has a progressive effect in our model.

The five steps described above comprise our particular approach to evaluating lifetime tax incidence. In general, our analysis emphasizes the estimation of lifetime incomes using panel data, the use of detailed consumption data to determine consumption patterns, and the general-equilibrium determination of tax burdens. In order to apply the lifetime approach in routine policy analysis, it is possible to simplify the procedure. Some of these possible simplifications will be discussed below.

III. A MEASURE OF LIFETIME INCOME

In order to discuss how lifetime income can differ from annual income, we begin with a case where the two measures are identical. If each person's income never changed over time, then annual income would accurately reflect permanent income. Each path in Figure 1 would be flat, and individuals would not change annual income categories. In this case, the poorest annual income category would include the same individuals as the poorest lifetime category.

The first difference between the two measures therefore arises from hump-shaped income profiles. Many studies confirm that incomes rise during early years, level off during later working years, and fall during retirement. This pattern puts young and old lifetime-rich individuals into low annual-income groups. The different groupings can affect incidence results.

A second difference can arise simply with income volatility. Self-employed individuals with a mid-range permanent income might be placed into a high-annual-income category, or a low-annual-income category, depending on the year taken for study. Employed workers subject to temporary lay-offs may experience similar fluctuations in annual incomes.

A third distinction is that, while the annual incidence of capital taxation depends on fixed capital endowments, the lifetime incidence of capital taxation depends on inheritances and on the

shape of the earnings profiles. If this profile is steeply peaked, the individual must save more for retirement and bear more burden of capital taxation. In our results below, we find that incidence depends not only on the height of the peak but on the timing of the peak. For example, we find that the earnings of middle income groups tend to peak later in life, so those groups do not accumulate as much savings relative to richer and poorer groups. The burden of capital taxation falls on those whose earnings peak early and who therefore save more for later.

A fourth difference is that the composition of lifetime income varies less than the composition of annual income. Differences in the capital-share of annual incomes that arise from the average amount of life-cycle savings are not relevant in the lifetime perspective. The only relevant differences in the composition of lifetime income must derive from bequests and inheritances, or from variations in the timing of earnings relative to consumption. Therefore taxes that change relative factor prices have less effect on the sources side. Similarly, all individuals progress from one set of consumption goods when young to another set of consumption goods when old. The composition of spending may still depend on income, in the lifetime perspective, but it does not depend on age. Therefore taxes that change relative goods prices have less effect on the uses side.

These considerations suggest that all distributional effects of taxes are likely to be muted in the lifetime context. The progressivity of the personal income tax places low tax rates not just on the lifetime-poor, but also on the lifetime-rich who are young. In addition, high personal taxes may be paid by lifetime-poor individuals who happen to be at the top of their earnings hump. On the other hand, a progressive annual tax structure generates heavier burdens on individuals with more humped lifetime income profiles, all else equal.

To estimate lifetime incomes, we require longitudinal data for many individuals over many years. This analysis has only recently become possible, because the University of Michigan's Panel Study of Income Dynamics (PSID) has been asking the same questions of the

same individuals now for over 18 years. From the PSID, we draw a sample of 500 households that includes 858 adult individuals with information on wages, taxes, transfers, and various demographic variables for the years 1970-87. We include heads and wives in our sample, and for simplicity in defining the lifetime of a "household", we exclude households whose marital status varied over the sample period. For heads and wives separately, we estimate the wage rate as a nonlinear function of age, so that for each individual in the sample we can: 1.) predict the wage rate for years that come after as well as before the sample period, 2.) multiply the actual or estimated gross-of-tax wage rate by a total number of hours per year (e.g. 4000) to get the value of the endowment, and 3.) calculate the present value of this endowment for each person. Thus our level of well-being is defined by potential earnings, including the value of leisure. These levels are used to classify households into twelve groups according to lifetime ability-to-pay, where we define a household's lifetime income to be the average of the head's and wife's (if any) lifetime incomes. We start with ten deciles, but we separate the poorest 2 percent from the next poorest 8 percent, and we separate the richest 2 percent from the next richest 8 percent.

Also, we are interested in the timing of income, because the shape of an individual's lifetime income profile determines the composition of annual income. Therefore we reestimate the nonlinear wage profile separately for each of the twelve groups. In addition, we require information on the time path of personal income taxes paid, and transfers received, in order to set up a consistent benchmark data set with a path of consumer spending out of total available net-of-tax income.

Thus our lifetime income and classification differs from previous studies such as Pechman and Okner (1974) because of hump-shaped earnings profiles, volatility in annual income, the timing of the peak in earnings, the exclusion of capital income, the use of a life-cycle model, the use of heads and spouses (but not entire households) as the unit of account, and the decision to include leisure in the total value of endowment. The next logical question, therefore, is

whether these issues really matter. How is lifetime classification different from the standard sort of classification?

We find that the annual income categories do not match up with the lifetime income categories for the same individuals. For each of our 858 individuals, we calculate annual income in 1984 for classification into annual income deciles. As it turns out, only 21.1 percent of these individuals are in the same annual and lifetime income deciles, and only 46.1 percent are within plus-or-minus one of the same decile. Most of the differences occur when someone who is lifetime-rich is very young or very old and earns low annual income. If we label the bottom 30 percent of the population the "poor" and the top 30 percent the "rich", we find that 13.8 percent of the annual poor are lifetime rich, and 2.6 percent of the annual rich are lifetime poor. We conclude that annual and lifetime income classifications are too different to assume that annual and lifetime tax incidence will be similar.

The final wage profiles for several of the lifetime income categories are illustrated in Figure 2. Each is the profile for a representative individual in the category. Groups 1 and 2 represent the lowest 2 percent and next lowest 8 percent of the lifetime income distribution, respectively. Group 4 represents individuals between the 20th and 30th percentiles, and Group 9 consists of those between the 70th and 80th percentiles. Group 12 is the top 2 percent of the population, and Group 11 is the next highest 8 percent.

From Figure 2 it is apparent that individuals in the higher lifetime income categories are characterized by wage profiles that are more peaked. The richest category (Group 12) appears to be significantly more peaked than the second richest (Group 11), which suggests that our split of the top decile is important. Because we use life-cycle consumption behavior in our model, the increasing peakedness implies that higher lifetime income categories will save more for retirement and will have higher ratios of capital to labor income. This ratio is key to the determination of the incidence of a consumption tax relative to capital taxes.

In Figure 2 we can also see that these wage profiles peak at different points in the life cycle. For the first few income groups, the peak age has a slight upward trend, as group 2 peaks at age 39, group 4 peaks at age 47, and group 9 peaks at age 64. Then, at higher levels of income, the peak wage years come earlier, at age 50 for group 11 and at 47 for group 12. Thus the middle-income groups, with the later earnings peaks, do not need to save as much for retirement. The lowest income groups and the highest income groups have earlier peaks, save more during life, and bear more burden of capital taxation.

IV. THE GENERAL EQUILIBRIUM MODEL

Our approach to the analysis of lifetime tax incidence uses a rather sophisticated general-equilibrium model to account for the behavioral effects and excess burdens caused by taxes. We want to capture important influences of taxes on diverse household choices about labor supply, savings, and the consumption of different commodities. We therefore assume utility maximization to find demands for commodities and supply of factors. We also want to capture effects of taxes on each producer's use of labor and capital, so we assume profit maximization to find demands for factors. We solve for general equilibrium prices in order to capture the net impact of taxes when these behaviors are considered simultaneously.

We assume that consumer decisions are made in stages. To begin, the individual calculates the present value of potential lifetime earnings. This endowment is supplemented by government transfers, reduced by taxes, discounted at the after-tax interest rate, and augmented by a fixed initial inheritance. For computational simplicity, we assume that the consumer expects the current interest rate to prevail in all future periods.

One part of this lifetime endowment must be saved for a bequest upon death. We avoid the many possible motivations for individual bequests, or the many ways in which taxes might affect the size of those bequests (see Bernheim, 1991). Instead, we are concerned that life-cycle

saving by itself can only explain about half of the observed capital stock (Kotlikoff and Summers, 1981). In our model, part of the capital stock is attributable to the fact that individuals receive exogenous inheritances and are then simply required to leave comparable bequests at the end of life. Incidence results depend on the differences in these exogenous inheritances among groups. To achieve balanced growth, each group must add some additional savings to their inheritance before they make their bequest.

The rest of the present value of income is available for spending. Decisions are made in stages, as depicted in Figure 3. At the first stage, the consumer chooses how much to spend each period. This choice depends on our assumption for the individual's rate of time preference (.005 in the central case) and the elasticity of substitution among time periods (0.5 in the central case). We later test the sensitivity of results to these parameters. The consumer's choice about how much to spend each period is also affected by changes in the net rate of return (which starts at .04 in the central case).

At the second stage, the consumer allocates one period's "spending" between leisure and other consumption goods. This choice depends on our assumption regarding another elasticity of substitution (0.5 in the central case). We allow individuals to "buy" more leisure at a price equal to the foregone net-of-tax wage, instead of buying other goods. This choice is affected by taxes, and it also depends on age. Individuals in this model never fully retire, but the weight on leisure increases with age after they reach 60, in a way that reflects actual choices.

In the third stage, individuals decide how to allocate current consumption spending among 17 particular goods (such as food, alcohol, tobacco, utilities, housing, etc.). This decision function takes a "Stone-Geary" form, which means that a consumer at a given age has to buy a set of 17 "minimum required purchase" amounts and then allocates remaining spending according to a set of 17 "marginal expenditure shares." These 34 parameters are estimated for each of 12 age categories using data from the Consumer Expenditure Survey, as described

thoroughly in Fullerton and Rogers (1993). This Stone-Geary framework has several important implications. By making a portion of spending non-discretionary, it reduces the sensitivity of total consumption and saving to the net rate of return. In addition, because discretionary income may be spent in proportions different from minimum requirements, actual purchase proportions depend on total income. Required spending is relatively high for housing and gasoline, while discretionary spending is relatively high for clothing, services, and recreation. Thus the rich and the poor buy different bundles, and bear different burdens on the uses side.⁷

In the fourth stage of our consumer's allocation process, in Figure 3, the expenditure on each consumer good is divided by fixed coefficients among components drawn from a list of producer industries. No real "decision" is made here, but this step allows us to match up consumption data using one definition of commodities with production data using a different definition. For example, expenditure on the consumer good "appliances" is composed of portions from metals and machinery, from transportation, and from the trade industry.

Then, in the fifth and final stage of the decision tree, the consumer takes the spending on each industry output and allocates it between the corporate sector and the noncorporate sector. We assume that the corporate output is not identical to the noncorporate output in the same industry. Hand-carved furniture, for example, is not the same as manufactured furniture. The consumer chooses the amount of each, using a weighing parameter based on initial observed

⁷This framework also allows us to use the same utility function for everyone in the model. In previous efforts, rich and poor individuals spend in different proportions because they have different preferences. But then the rich and the poor differ in fundamental characteristics and not just by the amount of income they receive. We feel that this assumption is very arbitrary: even if the poor were to receive additional income, they would still spend it as if they were poor, according to their unchanged proportions. It seems more natural, to us, that a poor person with more money would begin to behave like a rich person. That is, the primary distinction between rich and poor is the amount of income they receive. Therefore, in our model, everyone has the same preference parameters. The poor spend more on goods with high minimum required expenditures, because they are poor, and the rich spend more on goods with relatively high marginal expenditure shares.

corporate and noncorporate shares of production within each industry, and using another elasticity of substitution (5.0 in the central case). This specification allows us to capture the observed co-existence of both sectors within an industry, despite different tax treatments. If the outputs were identical, then a higher tax rate would drive one sector out of production. The degree of similarity is reflected in the elasticity of substitution. The other purpose of this specification is to capture ways in which changes in corporate taxes affect relative product prices and quantities demanded of the outputs of each sector.

We employ a similar decision tree, depicted in Figure 4, to model producer behavior in each sector of each industry. Each output is produced by many competitive firms in multi-stage production functions with constant returns to scale. Also, for computational simplicity, we assume no externalities, no adjustment costs, and no uncertainty.

In the first stage of production, output is composed of a fixed coefficient combination of value-added and intermediate inputs. Each of the 19 industries uses the outputs of all other industries, in fixed proportions. Thus we capture the effect of one product price upon another. In the second stage, value-added is a function of labor and "composite" capital. The weighing parameters are based on observed labor and capital in each industry, and the elasticity of substitution varies by industry (between .68 and .96, in the central case). Thus a tax on labor can induce the firm to use more capital instead, and vice versa. It also raises the cost of production, and thus output price, in any industry that uses a high proportion of the taxed factor.

In the third and final stage of production, composite capital is a function of five asset types -- equipment, structures, land, inventories, and intangibles. These types are defined by important tax differences such as the investment credit for equipment and the expensing of new intangible assets created through advertising or research and development. The weighting shares are again based on observed use of assets in each industry, and the response to tax differentials is again specified by an elasticity of substitution (1.5 in the central case).

Government in this model conducts several functions. It pays transfers to individuals according to the estimated lifetime transfer profiles discussed in the previous section. It produces an output for sale through an industry called "government enterprises," and it also produces a free public good through a composite combination of its use of labor, capital, and purchases of each private industry output. The weights in this combination are based on observed government purchases, and the elasticity of substitution is one. The level of this public good is held fixed in all simulations, as any tax change is accompanied by an adjustment that ensures equal-revenue yield. A final government function, of course, is to collect taxes.

Each tax instrument enters the model as a wedge between the producer's price and the consumer's price. The payroll tax, for example, applies an ad valorem rate to each producer's use of labor, so the gross-of-tax wage paid by the producer is higher than the net-of-tax wage received by the worker. Similarly, sales and excise taxes appear as ad valorem rates on each consumer good, so the gross-of-tax price paid by the consumer exceeds the net-of-tax price received by the seller.

The personal income tax is a little more complicated, in order to capture its progressive effect on tax burdens. The actual U.S. personal tax system imposes higher effective tax rates on higher incomes through a graduated rate structure with a changing marginal tax rate. For some purposes, one must calculate the effects of individual choices at each different possible marginal tax rate in order to determine utility-maximizing behavior. Our primary goal, however, is to measure the distributional effects of the tax. For this purpose, it is sufficient to use a set of linear tax functions that approximate the U.S. system with a negative intercept for each group and a single marginal tax rate (.3 in the central case). Although all individuals face the same marginal tax rate, average tax rates still increase with income due to the negative intercepts. We do not model the myriad exemptions and deductions. These simpler, linear tax functions can replicate the observed data on personal taxes actually paid by each group.

The state and local property tax and the U.S. federal corporate income tax raise the producer's gross-of-tax cost of capital, for each asset type, relative to the investor's net-of-tax rate of return. A Hall-Jorgenson (1967) formula shows how the cost of capital for each asset depends on the statutory corporate tax rate, depreciation allowances at historical cost, the rate at which inflation erodes those allowances, the rate of investment tax credit, and the required net rate of return for the firm. This required rate of return depends, in turn, on the going market rate and the personal taxation of interest, dividends, and capital gains. A similar cost of capital formula applies to the noncorporate sector. This treatment allows the producer's choice among assets to depend on relative tax rules, and the price of output in each industry to depend on the relative use of assets with different effective tax rates.

Other assumptions help to close the model in a way that accounts for all flows and that helps facilitate computation. We ignore international mobility of labor or capital, but allow for trade of industry outputs. Also, the value of imports must match the value of exports; the government's expenditures and transfer payments must match tax revenue; and, the value of personal savings must match the value of investment expenditures. Producer investment is not the result of firms' intertemporal optimization, but instead follows personal savings from consumers' optimization. The amount of personal savings is growing over time, because consumers' labor earnings are growing through population and technical change. On the steady state growth path, the capital stock grows at exactly the same rate as the effective labor stock.

Data for this model derive from many sources, adjusted to represent 1984 as the base year. In addition to the survey data used to estimate wage profiles and preference parameters, we use the National Income and Product Accounts for an input-output matrix, labor compensation by industry, government purchases, and international trade. These published data are combined with other unpublished data on capital allocations and inheritances.

For some parameters, such as the elasticities of substitution, we assume particular values.

For other parameters, such as the Stone-Geary preferences, we have econometric estimates. Finally, for remaining parameters, we "calibrate" from data on actual allocations. We use the demand functions, and all initial prices and observed quantities, to solve backwards for the value of the parameter that would make that quantity the desired one. This procedure establishes a "benchmark" equilibrium, with existing tax rules and prices, such that all consumers are buying the desired quantities and supplying the desired amounts of each factor, while producers are using their desired amounts of factors to produce the desired output.

Thus, using all these parameters together, we can solve for an equilibrium with unchanged tax rules that replicates the benchmark consistent data. This provides an important check on the solution procedure. From this benchmark, we can alter any particular tax rule and see how much more or less the consumers want to buy of each good. The solution algorithm then raises the price of any good in excess demand, and lowers the price of any good in excess supply, until it finds a set of prices where the quantity supplied equals the quantity demanded for every good and factor. It "simulates" the effect of the tax change, to calculate all new prices, quantities, and levels of consumer utility. Our measure of the change in tax burden is the "equivalent variation," the dollar value of the change in utility measured in terms of benchmark prices. We compare the lifetime equivalent variations to each category's lifetime income in order to determine relative lifetime tax burdens and the lifetime incidence of the various taxes, relative to the consumption- or wage-tax replacement.

V. RESULTS

As discussed in the introduction, the incidence of any additional tax within the general equilibrium model is defined relative to the tax which it replaces. We focus on the lifetime incidence of a general consumption tax (Tables 1 and 2), and a wage tax (Tables 3 and 4). Although both types of replacement taxes are intertemporally efficient, the consumption tax base

is larger than the wage tax base due to the presence of bequests. The timing of the two taxes also differs, because life-cycle behavior generates consumption paths that are smoother than the peaked wage paths. The two replacements therefore can produce very different transitional effects, but our results emphasize the long-run, steady-state incidence of taxes.

The distributional effects between rich and poor are measured by the gains to each lifetime income group in the steady state. The results for the consumption-tax replacement using our central parameterization are shown in Table 1. Overall gains in the steady state are shown at the bottom of each column in the table. This steady-state welfare gain is not a pure efficiency measure, because it includes redistribution to later generations from current generations who may lose when the tax change is imposed. Indeed, we often find gains to future steady-state generations and losses to older transitional generations. For an overall measure of "efficiency," at the bottom of the table, we simply take the present value of all equivalent variations for all generations, discounted at the net rate of return.⁸

Each table includes a column for each U.S. tax instrument being removed. We first discuss the consumption-tax replacement. In the first column of Table 1, the effects of replacing the personal income tax range from a 5 percent loss of income for the poorest group to a 9 percent gain for the richest group. The pattern is clearly regressive. The bottom of the column indicates that this replacement would yield an overall steady-state gain equal to 2.6 percent of income. This amount is also equal to 21 percent of the personal income tax revenue being replaced. The present-value gain is .7 percent of income (5 percent of revenue). The difference is caused by losses to transitional generations. Those who are near retirement at the time of the

⁸This procedure implicitly puts lower weight on later generations, but the discounting is necessary in order to obtain a finite sum for a sequence that is growing over time because of technical progress and population growth. In order to avoid having a measure of efficiency that would increase with the number of individuals, we discount by population growth to calculate this present value EV for the size of the population alive at the time of the change.

change would have paid low personal income taxes over the remainder of their lifetimes, but the consumption tax applies to all their remaining consumption.

The .7 percent efficiency gain might be considered substantial, especially since the personal income tax and the consumption tax both distort labor supply decisions. But progressivity makes the personal income tax more distorting. The personal marginal tax rate exceeds the total tax as a fraction of income, as is necessary for that average tax rate to rise with income. Distortions depend on the .30 marginal tax rate, while revenue depends on the average rate. Thus the personal income tax is relatively distorting per dollar of revenue. In contrast, the uniform consumption tax needs a rate of only .12 to collect the same revenue.

In the second column, a consumption tax is used to replace all existing U.S. sales and excise taxes. Contrary to initial intuition, this introduction of a consumption tax is clearly progressive. The explanation is that current sales and excise taxes are more regressive because they apply at nonuniform rates. Actual incidence depends on the pattern of tax rates and the pattern of goods consumed by the different lifetime income categories. Indeed, the point of our Stone-Geary specification is to allow for rich and poor households to purchase goods in different proportions. As it turns out, poorer households consume proportionately larger amounts of the high-taxed goods such as gasoline, alcohol, and tobacco. Poterba (1989) may be right that a uniform consumption tax would be proportional to lifetime income with no bequests, but the sales and excise taxes are not uniform. The rate structure introduces regressivity.⁹

The difference between the 1 percent steady-state welfare gain and the overall 0.18 percent efficiency gain indicates how much of the steady-state gain is attributable to effects in the transition.

⁹Both the existing sales and excise taxes and the replacement consumption tax exempt untaxed leisure, but they both apply to spending out of inheritances. Fullerton and Rogers (1993) show how these considerations affect the regressivity of all of these consumption taxes.

The third column of Table 1 considers the replacement of U.S. payroll taxes with a uniform consumption tax. Again the consumption tax replacement is quite progressive on a lifetime basis. And again this result conflicts with initial intuitions. In a simple life-cycle model with no bequests, each individual's present value of lifetime labor income exactly equals the present value of consumption. Therefore a uniform consumption tax is equivalent to a flat wage tax. Here the equivalence is broken by bequests which rise as a fraction of total endowment for higher lifetime-income groups. These high income individuals are required to leave the same size bequest as they receive, but they get to consume out of the extra capital income during their lifetimes. Therefore the consumption tax base exceeds the labor tax base for them. The simple consumption tax is more progressive than a labor income tax.

The switch from the payroll tax to consumption tax hurts the elderly, raises revenue from them, and allows gains to subsequent generations. These intergenerational gains and losses are shown graphically in Figure 5, for three of the tax change simulations. Lifetime effects of replacing the payroll tax on labor are represented by the bottom, solid line. Small net losses are incurred by those "born" 20 to 60 years ago (with chronological age 40 to 80). Net lifetime effects are positive and growing for those born later. The small bump for those born five years after the change is only an artifact of our five-year snapshots. Those born ten years after the change, and later, receive the steady-state gain equal to .50 percent of lifetime income. As shown in Table 1, the present value of these net effects is a mere .05 percent of the present value of all lifetime incomes. Gains to steady state generations are almost completely offset by losses to transitional generations, so the consumption tax is no more efficient than the labor income tax. Both distort primarily only labor supply decisions. The difference is purely distributional. The payroll tax is more regressive than the consumption tax.

In the fourth column of Table 1, if the consumption tax replaces current property taxes, the distribution of gains is U-shaped. This replacement raises the net rate of return and affects

owners of capital on the sources side. Because wage profiles peak later for middle income groups in Figure 2 above, they save less. Thus the ratio of capital income to labor income is approximately U-shaped across the lifetime income groups. The property tax repeal also reduces the cost of housing and thus helps low-income groups on the uses side.

The efficiency gain is only .17 percent of income, for two reasons. First, the property tax applies primarily to use of capital in the housing sector and thus helps offset some of the misallocation due to high corporate taxes on other uses of capital. Thus its repeal worsens intersectoral distortions. Second, the property tax is not a big tax. The last line shows that the efficiency gain is 5 percent of revenue, higher than for some other tax instruments.

Finally, in the fifth column, the consumption tax is used to replace the corporate income tax. Gains are flat or just slightly U-shaped. All groups gain from this replacement, because the corporate tax is distortionary. The efficiency gain is a very high 84 percent of revenue, since the 1984 corporate tax in this model misallocates resources without raising much revenue.

An apparent implication is to repeal the corporate tax. It could be entirely replaced by a small consumption tax, with gains to every income group. While this simulation does indicate such a Pareto improvement, it does not indicate the best possible reform. The consumption tax replacement might be costly to introduce, it might be difficult to administer, and it might not be applied to all goods at a uniform rate as in this simulation. A different approach might "fix" the corporate tax in order to collect more revenue with less distortion. Such an attempt was undertaken in the Tax Reform Act of 1986.

The last column of Table 1 shows the distributional effects of replacing the entire U.S. tax system with a uniform consumption tax. The poorest lifetime income group would lose 2 percent of lifetime income, and the richest group would gain 12 percent of income. Thus, current U.S. taxes are progressive relative to the consumption tax. The first five columns indicate that most of this overall progressivity is attributable to the personal income tax.

Overall, under the standard parameters, replacement of the entire U.S. tax system with a national VAT would generate an efficiency gain equal to 1.25 percent of lifetime incomes or 3.75 percent of tax revenues.

We now vary one of the key assumptions of the model, the intertemporal substitution elasticity. This parameter sets the degree to which consumers will switch between present consumption and future consumption, and it therefore helps determine the responsiveness of savings to relative prices like the net rate of return to capital. Responsive savers can "avoid" a tax on capital by saving less, which ultimately decreases the marginal product of labor, decreases the wage, and "shifts" the burden onto labor.

In the central-case simulations, we used .5 for this elasticity. This value might be considered on the high end of the estimates produced by econometric studies,¹⁰ but it helps generate a capital stock in our model that is close to that actually observed in the benchmark data. If we simply reduced this elasticity from .5 to .25, the initial steady-state capital stock would fall to unreasonably low levels (about 32.5 percent of the measured capital stock). Therefore we lower this elasticity and simultaneously lower the rate of time preference. In combination, these respecifications leave us with an initial capital stock close to that of the central case simulations.

Our "alternative" specification reduces the intertemporal substitution elasticity from .50 to .25, and it reduces the rate of time preference from .005 to -.005. As a theoretical matter, this rate of time preference may be either positive or negative. As an empirical matter, many readers may prefer a value that is even larger than the .005 we use in the standard case. The only way to achieve a higher rate of time preference with the same capital stock, however, is to raise the .5 value for the intertemporal substitution elasticity.

¹⁰See Hall (1988) and other studies cited in Auerbach and Kotlikoff (1987, pp. 50-51). Engen's (1992) estimates fall between .30 and .38 for this parameter.

Results are shown in Table 2. Comparing the present values in this table with those in Table 1, we see that these efficiency gains were larger under the standard parameters. In our alternative, replacement of the entire tax system with a VAT provides an efficiency gain of only .86 percent of lifetime incomes or 2.54 percent of tax revenues. This difference arises because the standard case uses a higher intertemporal elasticity. Thus, in the standard case, the savings distortions under current taxes are larger, and the efficiency gains from removing those distortions are higher. The contrast between efficiency effects in Tables 1 and 2 is especially clear for replacing taxes on personal income or on property, since these taxes generate the biggest intertemporal distortions.

Tables 3 and 4 show results for the wage tax replacement, in order to emphasize the possible differences between consumption taxation and wage taxation. In general, the intragenerational patterns of tax burdens under the wage-tax replacement are very similar to those under the consumption-tax replacement, with one exception. The uniform consumption tax looked progressive relative to current sales and excise taxes that place higher rates on goods purchased by low-income groups, but the wage tax in Tables 3 and 4 appears to provide U-shaped welfare gains. Low-income groups still gain from removal of regressive sales and excise taxes, but high-income groups also gain because the wage tax does not apply to spending out of their larger inheritances. Also, note in Tables 3 and 4 that the replacement of current payroll taxes with a uniform wage tax is not a neutral tax change, due to the regressivity of the social-security tax system and slightly different labor tax rates across industries.

The overall level of gain to the steady-state generation is lower under the wage-tax replacement, for two reasons. First, the smaller tax base of the wage tax necessitates a higher replacement tax rate. These tax rates are shown in Table 5. Second, the wage-tax replacement is relatively more burdensome to younger generations, because the consumption-tax places a lump-sum levy on the elderly during the transition. As indicated in the bottom two rows of

Tables 3 and 4, overall efficiency gains are lower under the wage-tax replacement.

VI. USING THE LIFETIME HORIZON IN POLICY ANALYSIS

Our lifetime analysis has certain qualities that are desirable from an academic standpoint: the accounting for general-equilibrium price effects, the use of utility-based measures which include excess burden as well as taxes paid, the measurement of lifetime income to classify rich versus poor, and finally, the measurement of lifetime burdens. These four qualities need not always go together. One could use a general equilibrium model and utility-based welfare measures without lifetime income or tax burdens. Also, one could look at annual tax burdens on groups classified by lifetime income. Our model incorporates all four. Although these qualities are attractive from an academic standpoint, they might be viewed as obstacles from a policy standpoint. In particular, the lifetime perspective has often been criticized as impractical.¹¹

The most fundamental criticism of the lifetime perspective is that it inappropriately presumes life-cycle behavior. Based on the life-cycle model of consumption, academic literature has often argued that lifetime income is a better indicator of a current ability to pay taxes. Moreover, the pure life-cycle model assumes perfect capital markets. It suggests that current consumption is proportional to lifetime income but independent of current income. Critics of lifetime measures point to a large literature with evidence that capital markets are far from perfect; liquidity constraints cause current consumption to track current income much more closely than the life-cycle model would predict. In the real world, then, high lifetime income does not necessarily allow high current consumption, and thus does not provide high current ability to pay taxes. We have argued in Fullerton and Rogers (1993), however, that lifetime

¹¹Many of these criticisms are discussed in the methodology pamphlet produced by the Joint Committee on Taxation (1993), pp. 32-33 and 83-86.

income should not be taken as a measure of current ability to pay taxes. It is simply a measure of lifetime ability to pay taxes. This interpretation applies even when the life-cycle model does not. It suggests that the lifetime perspective is a supplement to annual analyses -- a way of assessing the long-run distributional effects of taxes. Policymakers might worry that current taxes reflect current ability to pay and that lifetime taxes reflect lifetime ability to pay.

This criticism of the life-cycle model does pertain to our particular general-equilibrium model, however, because consumers maximize lifetime utility subject only to a lifetime budget constraint. Thus the consumption paths generated in our model are almost certainly smoother than real-world data would suggest. We use the life-cycle model because no other model provides a simple theoretical foundation while better tracking real-world observations. We modify the standard life-cycle model by adding bequests and heterogenous households, and we use the behavioral framework in order to simulate tax changes. However, analyses of lifetime tax incidence do not require the assumption of life-cycle behavior. One could use actual data on assets and consumption, perhaps merged with panel data on wage income, and then make incidence assumptions (as in Pechman and Okner, 1974, or Davies, St.-Hilaire, and Whalley, 1984). But such a model could not simulate behavioral changes in response to a tax change.

Another criticism of the lifetime perspective, as discussed by Barthold (1993), is that tax policies never last a lifetime. Indeed, the details of the tax code are modified very frequently. To analyze any particular reform, however, one cannot predict subsequent changes (or else those known future changes could be analyzed as part of the reform). The model is not used to predict the future, but to provide analytical insights and numerical magnitudes related to a particular policy. To address conceptual questions about that policy, we hold "all else equal." Besides, the reverse criticism could be levied against the annual perspective: major reforms last more than one year, so why look only at one-year effects?

Another frequent complaint about the lifetime perspective is that actual data on lifetime

incomes are simply not available. Meanwhile, the procedure we describe above to infer lifetime incomes is thought to be too cumbersome. While measuring lifetime income is obviously much more complicated than measuring annual income, the use of proxies for long-run income might be considered a practical compromise. For example, as mentioned earlier, the CBO has used one year's total expenditure as a supplemental measure of the ability to pay taxes. With some kind of life-cycle smoothing behavior, total expenditures in one year would reflect long-run income. For this reason, however, the use of the total-expenditure proxy is subject to the same criticisms as the life-cycle model. If capital markets are imperfect, then consumption levels may track annual income more closely than lifetime income. The total-expenditure proxy may then better reflect annual income than lifetime income. An alternative, but still relatively simple, proxy for lifetime income could be the level of education attained, or perhaps an estimated function of education, race, gender, and other age-invariant variables. This latter approach is investigated in Rogers (1994).

Finally, as suggested by the JCT (1993), computations generated by lifetime or infinite-horizon models can be very sensitive to the choice of parameters. In particular, numerical results are strongly dependent on intertemporal parameters such as the rate of time preference, discount rate, and interest elasticity of savings (or intertemporal elasticity of substitution). This criticism merely emphasizes the importance of conducting sensitivity analyses. Furthermore, annual-horizon models still generate results that are sensitive to choices of other substitution elasticities. They only avoid sensitivity to intertemporal parameters by ignoring intertemporal effects. Thus they are unable to shed light on the dynamic effects of policies.

In summary, lifetime tax incidence can play a valuable role in policymaking, as long as it is viewed as a supplement to, and not a replacement for, annual tax incidence calculations. The particular analysis described in this paper could easily be simplified in order to avoid general equilibrium price calculations and utility-based welfare measures. Alternatively, it could

be made more complicated in order to account for imperfect capital markets, liquidity constraints, endogenous bequests, or other phenomena. Either way, our purpose here is just to emphasize the basic distinction provided by the lifetime framework -- the measurement of lifetime tax burdens with respect to lifetime abilities to pay taxes.

VII. CONCLUSION

The lifetime tax incidence model described in this paper provides new insights into the long-run effects of consumption or wage taxation. Our simulation results highlight the tradeoffs between efficiency and equity in the debate about the choice between income-based and consumption-based taxes. For example, intertemporal efficiency is greatly improved by replacing the personal income tax with a consumption or wage tax, but such a switch is lifetime regressive in nature. A switch from either sales and excise taxes or payroll taxes to a uniform consumption tax is a progressive reform, but generates a much smaller efficiency gain. Our results suggest that the removal of the 1984 version of the corporate income tax provided the largest gain in relation to the small bit of revenue lost. For this tax switch, our distributional results are somewhat surprising. The corporate tax replacement would be slightly progressive in nature, because the corporate tax applied at higher effective rates to goods consumed by low-income families.

While the concept of lifetime tax incidence is widely accepted within the academic community, it has yet to be fully implemented by government tax-analysis agencies. We have attempted in this paper to respond to some of the criticisms of the lifetime approach, as voiced by government economists, but further efforts are needed to bridge the current gap between the rather stylized, academic analysis presented in this paper and the data-intensive types of studies undertaken by policy economists. Only then will the lifetime framework be routinely incorporated, along with the annual perspective, into real-world analyses of tax policies.

Figure 1
Two Lifetime Profiles




Figure 2

Wage Profiles for Selected Groups

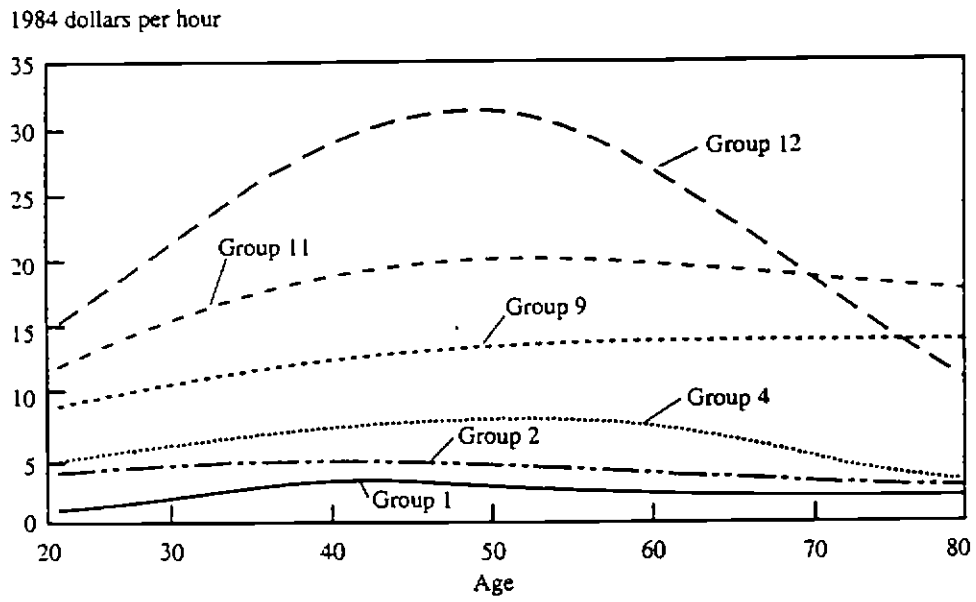


Figure 3

The Consumption Side of the Model

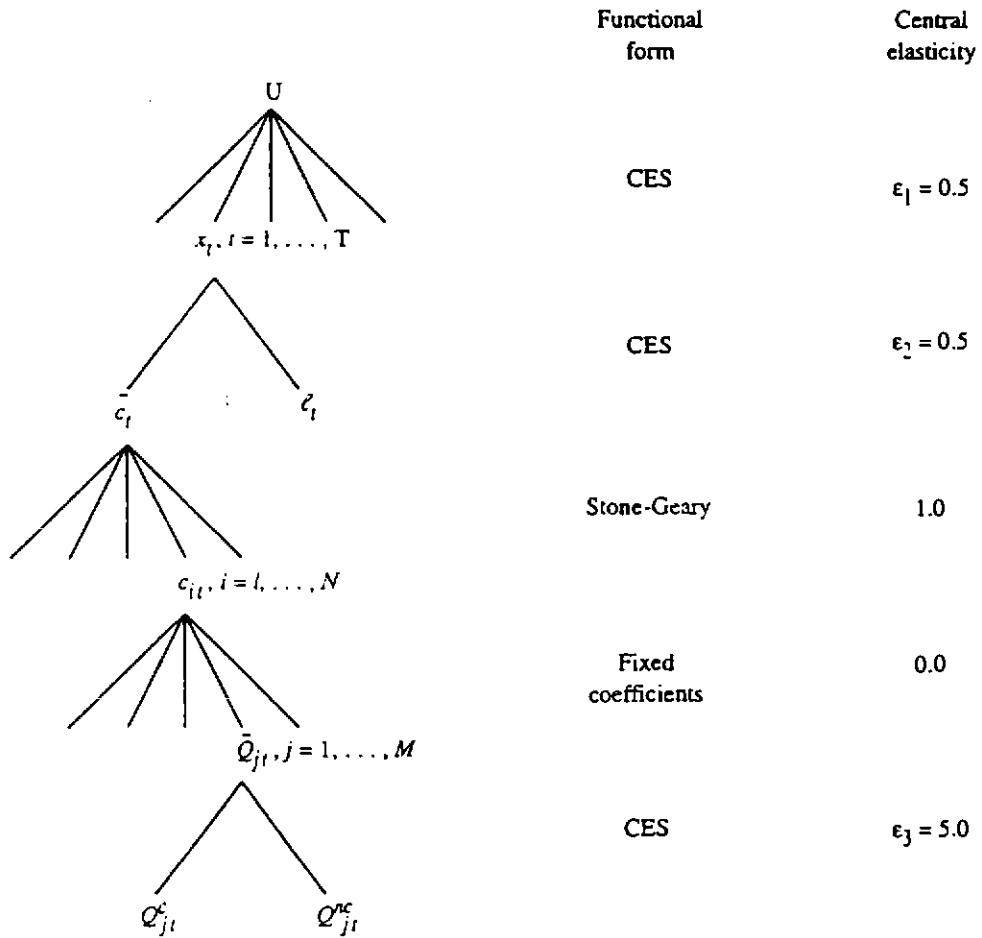


Figure 4

The Production Side of the Model

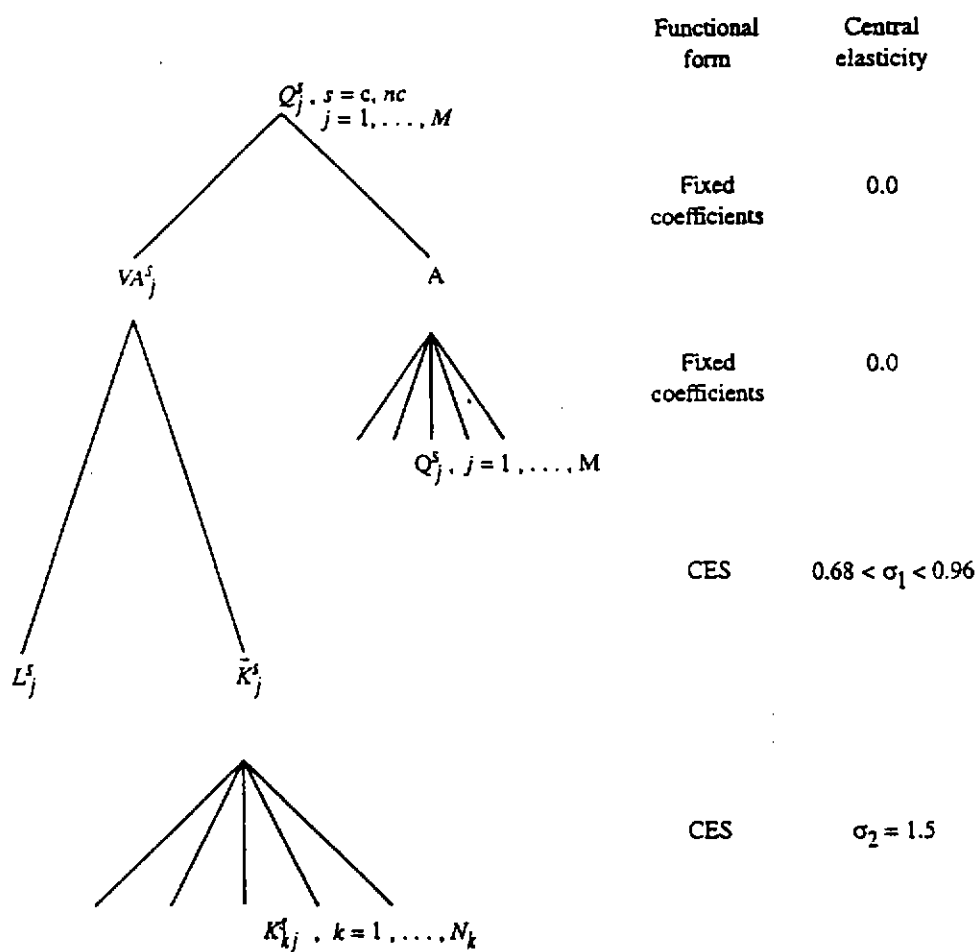


Figure 5

Gains and Losses to Each Generation
When a Consumption Tax Replaces Each U.S. Tax
(All Income Groups, Standard Parameters)

Equivalent variation as percent of lifetime income

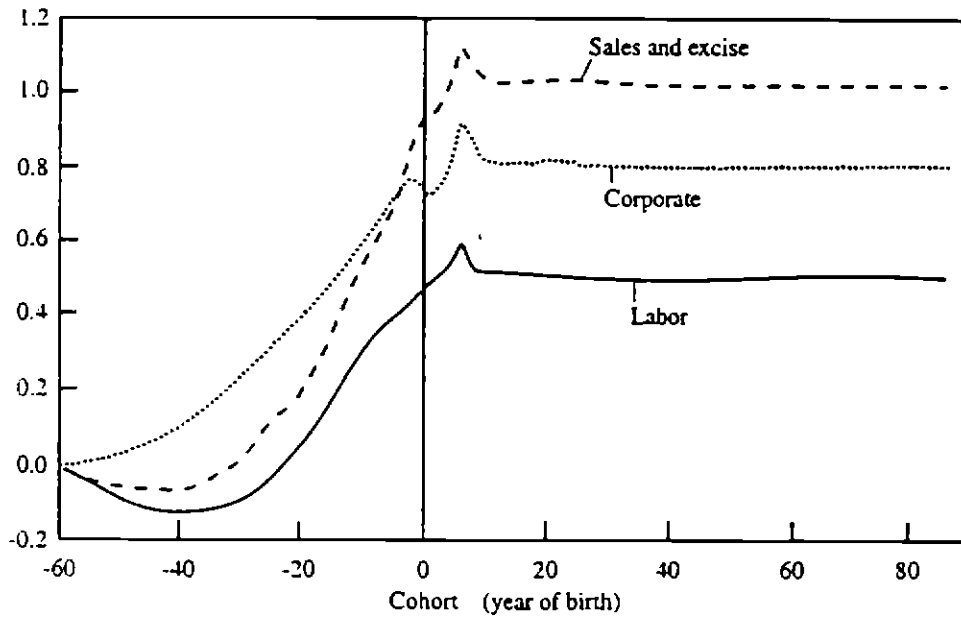


Table 1

Welfare Effects of a Uniform Consumption Tax, Replacing Each U.S. Tax, with Standard Parameters
 ($\epsilon_1 = .5, \delta = .005$)

EV as % of Lifetime Income, for Steady-State Generations

Lifetime Income Category	Personal Taxes	Sales & Excise Taxes	Payroll Taxes	Property Taxes	Corporate Taxes	Entire Tax System
1	-5.09	1.70	0.73	0.60	1.14	-2.00
2	-0.34	1.27	0.48	0.55	0.89	3.19
3	-0.70	1.22	0.54	0.21	0.78	2.48
4	0.71	1.16	0.55	0.33	0.81	3.94
5	2.53	1.19	0.64	0.09	0.78	5.66
6	0.66	1.11	0.61	0.19	0.77	3.74
7	2.22	1.04	0.53	0.59	0.85	5.64
8	2.43	1.08	0.65	0.14	0.76	5.47
9	3.01	1.06	0.67	0.17	0.77	6.09
10	3.04	0.98	0.55	0.16	0.74	5.83
11	4.19	0.72	0.16	0.73	0.83	6.81
12	9.11	0.60	0.01	1.08	0.93	11.87
All Twelve in Steady State:						
as % of Lifetime Income	2.57	1.02	0.50	0.36	0.80	5.58
as % of Revenue	20.88	10.32	8.86	11.58	314.63	17.62
All generations:						
as % of PV(Lifetime Income)	0.66	0.18	0.05	0.17	0.25	1.25
as % of PV(Revenue)	5.10	1.76	0.80	5.04	84.14	3.75

Table 2

Welfare Effects of a Uniform Consumption Tax, Replacing Each U.S. Tax, with Alternative Parameters
 ($\epsilon_1 = .25$, $\delta = -.005$)

EV as % of Lifetime Income, for Steady-State Generations

Lifetime Income Category	Personal Taxes	Sales & Excise Taxes	Payroll Taxes	Property Taxes	Corporate Taxes	Entire Tax System
1	-5.58	2.01	1.05	0.64	1.34	-2.85
2	-0.79	1.29	0.50	0.37	0.83	2.56
3	-1.38	1.27	0.61	-0.11	0.69	1.49
4	0.09	1.21	0.61	0.03	0.73	3.03
5	1.76	1.26	0.72	-0.31	0.66	4.53
6	-0.06	1.17	0.69	-0.17	0.67	2.71
7	1.80	1.06	0.56	0.41	0.80	5.07
8	1.62	1.14	0.73	-0.24	0.65	4.31
9	2.24	1.13	0.75	-0.21	0.67	4.99
10	2.28	1.04	0.63	-0.20	0.64	4.71
11	3.91	0.72	0.17	0.65	0.79	6.36
12	9.11	0.57	-0.02	1.16	0.93	11.82
All twelve in steady state:						
as % of Lifetime Income	1.96	1.06	0.55	0.09	0.72	4.71
as % of Revenue	15.83	10.84	9.94	2.74	214.22	14.75
All generations:						
as % of PV(Lifetime Incomes)	0.42	0.18	0.06	0.07	0.22	0.86
as % of PV(Revenues)	3.22	1.78	1.02	1.97	54.84	2.54

Table 3

Welfare Effects of a Uniform Wage Tax, Replacing Each U.S. Tax, with Standard Parameters
 ($\epsilon_1 = .5$, $\delta = .005$)

EV as % of Lifetime Income, for Steady-State Generations

Lifetime Income Category	Personal Taxes	Sales & Excise Taxes	Payroll Taxes	Property Taxes	Corporate Taxes	Entire Tax System
1	-6.21	1.19	0.39	0.36	1.12	-5.86
2	-1.47	0.62	0.08	0.27	0.87	-0.80
3	-1.96	0.39	0.06	-0.11	0.76	-1.97
4	-0.55	0.33	0.07	-0.00	0.79	-0.55
5	1.15	0.19	0.07	-0.28	0.75	0.80
6	-0.69	0.15	0.07	-0.17	0.74	-1.03
7	1.01	0.22	0.07	0.27	0.83	1.32
8	0.97	0.02	0.06	-0.24	0.73	0.36
9	1.56	-0.02	0.07	-0.22	0.75	0.96
10	1.81	0.10	0.06	-0.16	0.72	1.43
11	3.77	0.59	0.06	0.62	0.82	5.00
12	9.03	0.75	0.08	1.06	0.94	11.12
All Twelve in Steady State:						
as % of Lifetime Income	1.43	0.26	0.07	0.07	0.78	1.53
as % of Revenue	6.54	1.29	0.65	1.14	162.55	2.26
All generations:						
as % of PV(Lifetime Income)	0.60	0.14	0.02	0.14	0.25	0.93
as % of PV(Revenue)	2.59	0.66	0.19	2.26	44.14	1.31

Table 4

Welfare Effects of a Uniform Wage Tax, Replacing Each U.S. Tax, with Alternative Parameters
 $(\epsilon_1 = .25, \delta = -.005)$

EV as % of Lifetime Income, for Steady-State Generations

Lifetime Income Category	Personal Taxes	Sales & Excise Taxes	Payroll Taxes	Property Taxes	Corporate Taxes	Entire Tax System
1	-6.80	1.42	0.66	0.38	1.31	-6.78
2	-1.96	0.59	0.08	0.07	0.81	-1.42
3	-2.75	0.33	0.05	-0.46	0.66	-3.12
4	-1.28	0.27	0.06	-0.33	0.69	-1.64
5	0.26	0.11	0.06	-0.72	0.62	-0.55
6	-1.53	0.08	0.06	-0.56	0.63	-2.28
7	0.55	0.19	0.07	0.06	0.77	0.73
8	0.03	-0.05	0.05	-0.66	0.60	-1.05
9	0.65	-0.09	0.06	-0.64	0.62	-0.37
10	0.94	0.03	0.05	-0.56	0.60	0.16
11	3.54	0.59	0.06	0.56	0.79	5.01
12	9.19	0.80	0.08	1.17	0.95	11.75
All twelve in steady state:						
as % of Lifetime Income	0.77	0.22	0.06	-0.23	0.69	0.63
as % of Revenue	3.46	1.05	0.60	-3.72	105.67	0.92
All generations:						
as % of PV(Lifetime Incomes)	0.34	0.12	0.02	0.04	0.22	0.56
as % of PV(Revenues)	1.47	0.55	0.18	0.55	27.85	0.78

Table 5

Tax Rates for Each Replacement, with Standard Parameters
 ($\epsilon_1 = .5, \delta = .005$)

Removal Of:	Consumption Tax Rate		Wage Tax Rate	
	First Period	Steady State	First Period	Steady State
Personal Taxes	.1422	.1237	.1590	.1524
Sales & Excise Tax	.1104	.1073	.1469	.1457
Payroll Taxes	.0642	.0627	.0741	.0738
Property Taxes	.0384	.0328	.0460	.0414
Corporate Taxes	.0058	.0026	.0075	.0034
Entire Tax System	.4425	.3686	.4716	.4718

REFERENCES

Auerbach, Alan J. and Laurence J. Kotlikoff (1987), Dynamic Fiscal Policy, New York: Cambridge University Press.

Auerbach, Alan J., Laurence J. Kotlikoff, and Jonathan Skinner (1983), "The Efficiency Gains from Dynamic Tax Reform," International Economic Review 24, February, 81-100.

Ballard, Charles L., Don Fullerton, John B. Shoven, and John Whalley (1985), A General Equilibrium Model for Tax Policy Evaluation, Chicago: University of Chicago Press (for National Bureau of Economic Research).

Barthold, Thomas A. (1993), "How Should We Measure Distribution," National Tax Journal 46, September, 291-99.

Bernheim, B. Douglas (1991), "How Strong Are Bequest Motives? Evidence Based on Estimates of the Demand for Life Insurance and Annuities," Journal of Political Economy 99, October, 899-927.

Davies, James B., France St-Hilaire, and John Whalley (1984), "Some Calculations of Lifetime Tax Incidence," American Economic Review 74, September, 633-49.

Engen, Eric M. (1992), "Precautionary Saving, Consumption, and Taxation in a Life-Cycle Model with Stochastic Earnings and Mortality Risk," unpublished PhD dissertation, University of Virginia.

Friedman, Milton (1957), A Theory of the Consumption Function, Princeton: Princeton University Press.

Fullerton, Don and Diane Lim Rogers (1993), Who Bears the Lifetime Tax Burden?, Washington DC: The Brookings Institution.

Hall, Robert E. (1988), "Intertemporal Substitution in Consumption," Journal of Political Economy 96, April, 339-57.

Hall, Robert E. and Dale W. Jorgenson (1967), "Tax Policy and Investment Behavior," American Economic Review 57, June, 391-414.

Kotlikoff, Laurence J. and Lawrence H. Summers (1981), "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," Journal of Political Economy 89, August, 706-32.

Lyon, Andrew B. and Robert M. Schwab (1990), "Consumption Taxes in a Life-Cycle Framework: Are Sin Taxes Regressive?", mimeo, University of Maryland.

Pechman, Joseph, A. and Benjamin A. Okner (1974), Who Bears the Tax Burden?, Washington, D.C.: Brookings Institution.

Poterba, James M. (1989), "Lifetime Incidence and the Distributional Burden of Excise Taxes," American Economic Review 79, May, 325-30.

Rogers, Diane Lim (1994), "Measuring the Distributional Effects of Corrective Taxation," mimeo, Pennsylvania State University and U.S. Congressional Budget Office.

Summers, Lawrence H. (1981), "Capital Taxation and Accumulation in a Life Cycle Growth Model," American Economic Review 71, September, 533-44.

U.S. Joint Committee on Taxation (1993), "Methodology and Issues in Measuring Changes in the Distribution of Tax Burdens," Joint Committee Print, Washington DC: U.S. Government Printing Office, June.