

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Tax Policy and the Economy, volume 13

Volume Author/Editor: James M. Poterba, editor

Volume Publisher: MIT Press

Volume ISBN: 0-262-66150-0

Volume URL: <http://www.nber.org/books/pote99-2>

Publication Date: January 1999

Chapter Title: Distributional Impacts of Proposed Changes to the Social Security System

Chapter Author: Julia Lynn Coronado, Don Fullerton, Thomas Glass

Chapter URL: <http://www.nber.org/chapters/c10924>

Chapter pages in book: (p. 149 - 186)

DISTRIBUTIONAL IMPACTS OF PROPOSED CHANGES TO THE SOCIAL SECURITY SYSTEM

Julia Lynn Coronado

Board of Governors of the Federal Reserve System

Don Fullerton

University of Texas at Austin and NBER

Thomas Glass

University of Texas at Austin

EXECUTIVE SUMMARY

In this paper we assess the degree to which the current social security system redistributes income from rich to poor. We then estimate the impact of various proposed changes to social security on the overall redistributive effect of the system. Our analysis takes a steady-state approach in which we assume participants work their entire lives and retire under a given system. Redistribution is measured on a lifetime basis using estimated earnings profiles for a sample of people taken from

This paper was prepared for the National Bureau of Economic Research (NBER) Conference on "Tax Policy and the Economy" October 20, 1998, in Washington, DC. We are grateful for helpful comments from Gib Metcalf, Ed Olsen, Jim Poterba, and Jon Skinner. 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 Federal Reserve Board or the National Bureau of Economic Research.

the PSID. We allow for differential mortality, not only by gender and race, but also by lifetime income. Our results indicate that the current social security system redistributes less than is generally perceived, mainly because people with higher lifetime income live longer and therefore draw benefits longer. Remaining progressivity is reduced and even reversed by an increase in the assumed discount rate, since regressive taxes become more important relative to later progressive benefits. We find that many of the proposed changes to social security have surprisingly little effect on the redistribution inherent in the system.

1. INTRODUCTION

The goal of this paper is to analyze ways in which the current social security system and some proposed reforms redistribute between high- and low-income groups, defined on a lifetime basis. Rather than look at redistributions between age cohorts, our analysis focuses exclusively on intragenerational redistribution in a steady state. We assume that all working years and retirement years come under a single social security system. Thus we assess long-run redistributive effects of the current system and of several reforms. Within this steady-state context, we look at the system's lifetime redistribution across groups defined by income, gender, and marital status. We allow for heterogeneity within each income group, as specific features of the social security system differentially affect groups with different proportions of individuals who are single or married, are male or female, work continuously or sporadically, and have different mortality rates.

Our analysis proceeds in five stages. In the first stage, we use 22 years of wage rates from the Panel Study of Income Dynamics (PSID) to estimate wage-rate profiles for different kinds of individuals (household heads, full-time working wives, and part-time working wives). The estimated coefficients are used to project each individual's wage rates before and after the sample period, so that each individual has a complete wage profile from age 22 to 66. The wage rate for each year is multiplied by a total time endowment to calculate potential earnings, and the present value of this endowment is used to categorize individuals into quintiles from rich to poor.

In the second stage, for each quintile, actual earnings are used to estimate earnings profiles. We again use the coefficients to project out-of-sample earnings for each individual, so that each member of our sample has a complete lifetime earnings history. In the third stage, we derive income-differentiated mortality rates for each group. Then, in the fourth stage, we use the constructed earnings histories and mortality

probabilities to calculate each individual's expected lifetime social security taxes and benefits. In the final stage, we add over the individuals in each quintile to get the net impact of social security on each group. We also calculate the redistributive effects of four proposed reforms.

Using actual earnings data is one of the important innovations of our paper. As noted below, previous studies use stylized groups, or smoothly-estimated profiles for each group. In contrast, the use of actual earnings data allows us to incorporate differential effects of human capital investment, illnesses, child rearing, and other events that affect earnings and that may lead individuals to enter and exit the labor force. We also give special attention to differential mortality rates by gender, race, and lifetime income.

Distributional effects of the current system also represent the effects of a major reform, namely, repeal of social security. In addition, we calculate effects of four smaller reforms designed to reduce the current social security deficit by the same amount: eliminating the provision for dropping certain low-earnings years from the benefit calculation; increasing the age of retirement; increasing the tax rate; and decreasing benefits.

We find that: (1) overall, the social security system is progressive; (2) allowing for income-differentiated mortality substantially reduces measured progressivity; (3) increasing the assumed discount rate can eliminate remaining progressivity; (4) the four reforms we study are somewhat regressive; and (5) income-differentiated mortality lessens the regressive nature of the proposed reforms.

2. OUR APPROACH RELATIVE TO PREVIOUS WORK

To clarify how and why our approach differs from existing literature, consider the two illustrative lifetime wage profiles in Figure 1. The relatively-poor person's wage increases with age through points *A*, *B*, and *C*, and then falls to point *D* at retirement. The rich person's profile is higher (through points *E*, *F*, *G*, and *H*). In this context, the social security system may take taxes from both types of persons during working years, and provide benefits to both when retired. We wish to measure how much of this money is transferred from the high-lifetime-income person to the low-lifetime-income person, rather than just transferred from working years to retirement years within the same group.

Initial tax incidence studies like Pechman and Okner (1974) used groupings based on annual income. This type of study finds that the social security system is progressive, but it aggregates unlike individuals. The richest group includes only those at point *G*, the next group includes those at points *F* and *C*, the following group includes those at




FIGURE 1. *Two Lifetime Profiles*

E, B, and H, and the poorest group includes very young and old individuals at points A and D.

Some later studies like Auerbach and Kotlikoff (1987) include lifetime profiles and lifetime decisionmaking, in order to find how social security redistributes between young and old. However, the youngest group aggregates individuals at points A and E, the next group includes B and F, and the oldest group includes D and H. This type of study also does not distinguish between the two lifetime income groups in Figure 1.

Although much work has focused on intergenerational effects of the social security system, considerable work has also looked at intra-generational redistribution—using arbitrary levels of income for different groups. For example, Hurd and Shoven (1985) and Boskin, Kotlikoff,

Puffert, and Shoven (1987) each use three groups (e.g. median income, half the median, and five times the median).¹ As in Figure 1, the shape of the earnings profile does not differ between the three groups.

The approach of using arbitrarily-set income levels has tremendous computational appeal. However, the calculation of social security benefits does not depend just on the level of lifetime earnings. Recent years often get more weight, and some years with zero earnings can be dropped from the calculation. Thus the benefits received by each group depend on the *shape* of the earnings profile and the *variance* from one year to the next. For these reasons, we estimate a nonlinear profile separately for each group. We retain actual earnings data from the sample period and use actual and constructed years of data with zero earnings. Each group has different proportions of individuals with different numbers of zero-earnings years that can be dropped from the benefit calculations (as in Williams, 1998).

Some studies have used actual social security records to look at issues of redistribution. Burkhauser and Warlick (1981) and Hurd and Shoven (1985) use extracts from social security records, while Duggan, Gillingham, and Greenlees (1993) use records for more than 32,000 workers from the Continuous Work History Sample of social security records. While using social security records would better identify social security earnings histories, two important elements are missing from the available extracts. First, the observed amount of earnings is generally capped at the annual social security wage cap. Yet only data with wage rates above the cap can capture the regressivity of social security taxes that exempt higher wages. Second, and equally important, records for individuals are not linked with records of spouses.

Fullerton and Rogers (1993) also estimate profiles separately for 12 different lifetime income groups, and use them to calculate the incidences of various taxes, but they do not look at social security benefits. More recently, Altig et al. (1997) employ the same 12 lifetime income groups in their model of tax incidence, and Kotlikoff, Smetters, and Walliser (1998) do use that model to look at social security. These computational general equilibrium models can calculate the effects of social security reforms on factor returns in each period, but each of the 12 groups is assumed to be homogeneous. Since everyone in a group must work the average amount for that group these general equilibrium mod-

¹ Panis and Lillard (1996) set the low group at full-time minimum-wage earnings, the middle group at the social security average earnings, and the high group at the social security tax wage cap. Similar procedures are followed by Myers and Schobel (1983), Steuerle and Bakija (1994), and Garrett (1995).

els cannot incorporate heterogeneity such as the fraction in each group who have zero earnings.

For these reasons, in this paper we do not attempt to build a general equilibrium model. The point of this paper is to make use of actual data on diverse individuals within each lifetime income group. We can thus use the fact that each group has a different proportion of individuals with zero-earnings years, a different proportion who qualify for spousal benefits, and a different proportion who receive fewer benefits because they die earlier. In this way, we can look at distributional effects of specific elements of the social security system.²

The literature on distributional effects of specific elements of the social security system is small. Flowers and Horowitz (1993) look at the spousal benefit, whereby low-earner spouses can draw the greater of their own computed benefit and one-half the higher-earning spouse's benefit. They demonstrate that the spousal benefit calculation is progressive compared to an own-benefit calculation. This result is driven by their finding that higher-income families consist of spouses with more equal earnings (lower-income couples have more disparate earnings). Panis and Lillard (1996) use a low-medium-high income structure to examine three basic reforms: increasing retirement age, increasing payroll taxes, and decreasing benefits.

A few studies introduce income-differentiated mortality into analysis of the social security system. Rofman (1993) uses a data set that matches demographic information from the Current Population Survey with social security information on earnings, benefits, and mortality. However, Duleep (1986) reports that mortality information is severely underreported in the social security records, especially for working-age individuals and for minorities. Garrett (1995) uses mortality estimates from a literature search, and Panis and Lillard (1996) extract mortality information from the PSID. Since high-income people live longer, several studies show that allowing for income-differentiated mortality seriously dampens the progressivity of social security (e.g. Steuerle and Bakija, 1994; Duggan, Gillingham, and Greenlees, 1995; and Panis and Lillard, 1996).

Finally, Caldwell et al. (this volume) use *CORSIM*, a large micro-simulation model, to construct lifetime earnings for many individuals. This model starts with the 1960 Census Public-Use Microdata Sample and uses estimated transition probabilities to grow the sample in one-

² By concentrating on dollar flows, however, we miss the effect of this social insurance program on the utility of risk-averse individuals. The benefits of risk reduction may be larger for low- or high-income individuals. Lee, McClellan, and Skinner (this volume) calculate such effects for Medicare.

year intervals. For each person, they simulate the next year's income and work status. Thus, as in our study, they capture differences in race, gender, the number of zero-earnings years, differential mortality, and wage rates above the cap. They focus primarily on intergenerational redistributions, finding that while early generations received a good rate of return, postwar generations receive smaller and even negative rates of return.

3. WAGE PROFILES, QUINTILES, AND EARNINGS HISTORIES

The data and methodology used to obtain our lifetime wage profiles and earnings profiles are summarized in the Appendix.³ We use the PSID for the years 1968 to 1989, which gives us 22 years of actual earnings data for a sample of the population. Section A.1 discusses the selection of our sample, consisting of 1,082 heads and 696 wives. In the PSID, if a household contains a married couple, the husband is automatically designated as the household head. Thus, most heads in our sample are male. Of the 386 single heads of household, 118 are female and 268 are male.

In the first stage, we estimate wage profiles in order to calculate the present value of potential earnings and to categorize members of our sample into lifetime income quintiles. The goal of this stage is to divide the heads and wives into lifetime income groups that identify those who are rich or poor, using a broad measure of economic welfare. To include the value of leisure and the value of home production, we use potential earnings rather than actual earnings. We assume that each individual could work a maximum of 80 hours per week (4,000 hours per year). Then, in order to avoid the fluctuations of annual income, we classify individuals according to lifetime potential earnings. In addition, because of the difficulties in defining the lifetime of a household, we categorize people on an individual basis.

In the second stage, to look at the effect of social security, we need actual earnings information from ages 22 to 66. For this reason, we then use our PSID data to estimate earnings profiles by quintile and gender, and we use the coefficients for each individual to construct earnings for out-of-sample years.

³ A full description appears in Coronado, Fullerton, and Glass (1998). That paper develops the model and uses it to evaluate earlier calculations of progressivity of the existing social security system. The current paper does not evaluate earlier calculations for the existing social security system, but instead evaluates proposed reforms.

3.1 *Lifetime Income*

We first estimate potential lifetime income. To begin, we divide each year's earnings by hours worked to calculate the annual wage rate separately for each head and wife. We then use this wage rate and multiply it in each year by 4,000 hours to represent the year's labor endowment. This product represents the potential earnings of the individual and therefore serves as a measure of his or her material well-being. Using this endowment allows us to abstract from the actual labor-leisure choice, since someone who chooses to consume more leisure might be just as well off as someone who decides to work more and consume less leisure. Using potential income also avoids the distortion introduced by the fact that home production does not show up in the data under hours worked. The wage rate is a measure of earning power that reflects experience, talent, and education.

To construct wage rates for every year of our sample members' working lives, we use the PSID data to estimate log wage profiles. We estimate separate log wage regressions for heads and wives. For the heads of household we take all positive observations for wages, which gives us 19,130 observations on 1,082 heads. We regress the log of the wage rate on an individual fixed effect and other variables like age, age squared, and age cubed. We thus estimate a shape of the wage profile like those in Figure 1. Because we have a fixed effect for each individual, we cannot use variables that do not vary over time (like race or gender). However, we do include age *interacted* with education, race, and gender. The results of this regression are shown in Table 1. Using the resulting fixed effects and coefficients, we then fill in missing observations during the sample period, and we fill in observations outside the sample period. Thus, for each household head, we have a wage rate for every year of their entire economic life from age 22 to 66.

To capture potential earnings for each wife, we want a wage rate for each year whether or not she worked. Non-working wives do engage in household production, and assigning them a zero wage may incorrectly place them in a low lifetime income group. The data indicate that wives fall into three basic categories: wives who work on a regular basis, wives who work in only a few of the sample years, and wives who do not work at all. For those wives who average more than 750 hours of work per year, we have 5,413 observations on 307 women. For those who work occasionally, but average less than 750 hours, we have 2,292 observations on 296 wives. A third group includes 93 wives who do not work at all throughout the sample.

For each of the two groups of working women, we take all positive

TABLE 1
Log Wage Regression for Heads of Household

Independent variable	Coefficient	<i>t</i> -statistic
Age	0.1343	6.26
Age ²	-0.003313	-8.53
Age ³	0.000026	9.55
Age × education	0.003669	4.87
Age ² × education	-0.0000326	-4.52
Age × female	-0.0239	-1.89
Age ² × female	0.000306	2.11
Age × white	0.0167	1.32
Age ² × white	-0.000240	-1.67
Individuals		1,082
Observations		19,130
Adjusted R ²		0.57

observations and regress the log of the wage rate on an individual fixed effect and variables for age and the interaction between age and education.⁴ The results of these regressions for the two groups of women can be found in our previous paper (Coronado, Fullerton, and Glass, 1998). We again use the estimated fixed effects and coefficients to fill in missing observations within the sample, and to simulate observations outside the sample, so that each wife has a complete wage profile for ages 22–66. For the 93 women who did not work at all during the sample, we assign them the median fixed effect for the women who averaged less than 750 hours of work annually. We then use the coefficients from that group's log wage regression to fill in the entire profile of potential hourly wages.

Once we have a complete wage profile for each of our heads and wives for ages 22–66, we calculate individual gross lifetime income as:

$$LI = \sum_{t=1}^{45} \frac{P_t(w_t \times 4,000)}{(1+r)^{t-1}},$$

where t indexes the 45 years in the individual's economic lifetime relevant for social security (ages 22 to 66), where the individual could work a maximum of 80 hours per week for 50 weeks per year, and where P_t is the individual's probability of survival to age t . We use two different values for r , the discount rate.

As couples generally pool their resources, it would be inappropriate to

⁴ The PSID does not have a race variable for the wives in the sample.

place husbands and wives individually into separate lifetime income quintiles. Thus the next step is to use a household equivalence scale that accounts for the average net economies and diseconomies of scale in the different categories of household consumption. Jorgenson and Slesnick (1987) find that this adjustment means taking the sum of the couple's incomes and dividing by 1.934, so the implication is net economies of scale.⁵ Most importantly, this use of equivalence scales is designed to guarantee that a husband and wife are always placed in the same lifetime income quintile, regardless of their separate incomes.

We can now deal with all members of our sample as individuals and categorize them into five lifetime income groups. The first quintile has the lowest income, and the fifth has the highest income.

3.2 Earnings Histories

In the first stage described above, we divide individuals into lifetime income quintiles based on a wage profile and the implied present value of their labor endowment. Individuals are thus separated into categories based on what we feel to be an appropriate measure of economic well-being. Because the costs and benefits of social security are based on actual income, however, we also need profiles of actual earnings. In this stage we estimate earnings regressions instead of wage regressions. For each quintile, using our data from the PSID, section A.2 describes how we estimate separate earnings regressions for heads, working wives, and occasional working wives. In this stage we use both positive and zero earnings observations. We then use the results of these Tobit regressions to simulate earnings for out-of-sample years, and we construct a complete profile of "actual" earnings for each individual from ages 22 to 66. Our methodology allows for the simulation of years with zero earnings. We then use these profiles in our analysis of the distributional effects of social security.

We estimate our earnings regressions using maximum likelihood separately by quintile for heads, habitual working wives, and part-time working wives. We thus estimate a total of 15 regressions (reported in Coronado, Fullerton, and Glass, 1998). For each regression for the heads, we begin with independent variables for age, age squared, age cubed, education, education squared, gender, race, the product of age and education, the product of age and gender, and the product of age and race. We then eliminate the variables that are insignificant. For the regressions for

⁵ For all of our couples, we use the equivalence scale estimated by Jorgensen and Slesnick (1987) for an urban household in the northeastern U.S. headed by a white person aged 35–44.

wives, we begin by including age, age squared, age cubed, education, education squared, and the product of age and education. We again eliminate the variables that were insignificant.

We next use the estimated coefficients from our earnings regressions to simulate earnings observations for the out-of-sample years for all individuals in our sample, so that each person has an earnings profile for ages 22 to 66. Unlike the first stage, we do not also use these coefficients to fill in missing or zero earnings observations during the sample period. This is because we are interested in actual earnings, and years spent out of the labor force are relevant for calculating the costs and benefits of social security for an individual. In fact, we also calculate a representative number of zero-earnings years in the simulated out-of-sample portions of each earnings profile.

Combining the actual observations with the simulated observations for each individual yields a complete earnings profile for ages 22 to 66. We next proceed to use these profiles to analyze the distributional effect of social security. The advantage of using these estimated profiles is that we can allow for entry and exit from the labor force. These events are relevant when evaluating the redistributive effect of social security, because benefits are based on earnings histories and allow for a certain number of years to be dropped before making average-wage calculations. Another advantage of using actual data to analyze the redistributive effect of social security is that we have a demographically-diverse sample. This diversity affects our analysis in that different demographic groups have different mortality rates. These differences turn out to be an important issue in analyzing social security, as described below.

4. INCOME-DIFFERENTIATED MORTALITY

Standard mortality tables extend only to age 85 and are differentiated only by sex and race. As described in section A.3, we extend these data in three ways. First, we describe assumptions necessary to extend the tables to age 99. Second, since individuals with low incomes have higher mortality rates than the population as a whole, we modify the standard tables by using available information on mortality differentiated by *annual* income. Third, we then use that information to construct mortality tables that are differentiated among our *lifetime* income quintiles. In later sections we use these tables to compute expected values of social security taxes and benefits.

Standard mortality tables are provided in *Vital Statistics of the United States* (U.S. Department of Health and Human Services, 1993). These tables show the number (out of 100,000) who remain alive, for each age

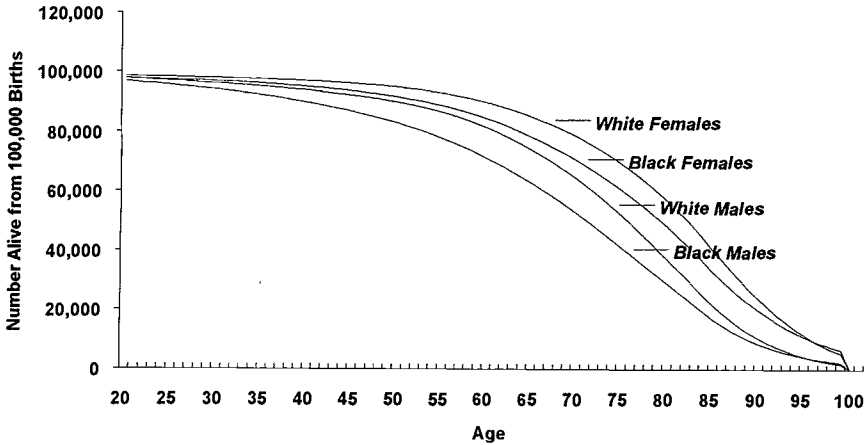


FIGURE 2. *Number Still Alive as a Function of Age*

Source: *Vital Statistics of the U.S.—1989*, expanded.

up to 85. Some prior studies use a simple procedure in which they compute normal life expectancy at each age and then assume that the individual will be alive exactly that long and will die at the date of life expectancy.⁶ Instead, we use the probability of remaining alive at each age. Based on standard mortality tables, a hypothetical 22-year-old white male has probabilities of survival to age 23 of 99.83 percent, survival to age 65 of 75.82 percent and survival to age 85 of 22.34 percent. We multiply the tax that would be due or the benefit that would be received at each age by the probability of attaining that age, and then calculate the rate of return on these *expected* cash flows. Because all outflows (taxes) occur in the early years and all inflows (benefits) occur in the later years, this method will differ considerably from the simpler procedure just described.

The National Center for Health Statistics obtains death certificates from all U.S. states and constructs four "current life tables" (for white males, white females, non-white males, and non-white females).⁷ For

⁶ For example, a white male aged 22 has a life expectancy of 51.1 years and can therefore expect to attain age 73.1. This simple procedure would determine the amount of taxes paid and benefits received through age 73.1 and then calculate the rate of return that equates these present values.

⁷ "Thus, for example, a current life table for 1989 assumes a hypothetical cohort subject throughout its lifetime to the age-specific death rates prevailing for the actual population in 1989. The current life table may thus be characterized as rendering a 'snapshot' of current mortality experience, and shows the long-range implications of a set of age-specific death rates that prevailed in a given year." (*Vital Statistics of the United States—1989*.)

TABLE 2
Ratio of Observed Deaths to Expected Deaths (O/E) for Each Race-Sex Group—Ages 25–34

Annual family income (\$)	Number <i>n</i>	Percentile	O/E			
			White male	White fem.	Other male	Other fem.
<5,000	11,670	6.31	1.68	1.51	1.54	1.86
5,000–9,999	22,085	18.25	1.20	0.97	0.81	1.01
10,000–14,999	33,331	36.27	1.28	1.17	1.36	1.01
15,000–19,999	32,231	53.70	1.12	0.76	0.71	0.84
20,000–24,999	30,729	70.31	0.80	0.97	0.92	0.36
25,000–49,999	48,375	96.47	0.73	0.94	0.72	0.44
>49,999	6,529	100.00	0.61	1.15	0.72	0.44
Totals	184,950		<i>n</i> = 81,461	<i>n</i> = 85,047	<i>n</i> = 7,752	<i>n</i> = 10,690

Source: Rogot, Sorlie, Johnson, and Schmitt (1992, Table 7). The "expected" number of deaths is based on the overall death rate within the age-sex-race category, not differentiated by income, while "observed" deaths are the actual deaths in each income group.

100,000 individuals alive at age 0, the table shows the number surviving at each age 1 through 85. Since 31 percent of the population is still alive at age 85, section A.3 describes how we extend the tables through age 99. These expanded mortality tables allow us to weight tax payments and benefits by the probability of dying in each year from age 22 to 99. Figure 2 shows the extended mortality for the four race-sex groups.

Many studies have noted that mortality rates for the poor are larger than average. *A Mortality Study of 1.3 Million Persons* (Rogot, Sorlie, Johnson, and Schmitt, 1992) provides a rich source of data on this effect. We use their information on observed deaths, *O*, and the number of deaths that would be expected if all income groups had the same mortality rate, *E*, and apply the *O/E* ratios to each cell in the extended mortality tables. Results for 25–34-year-olds are shown in Table 2, but we derive similar tables for each age group.⁸

Among white males, Table 2 shows that those in the poorest annual-income group die at a rate that is 168 percent of the average for their age group, while those in the richest annual-income group die at a rate that is only 61 percent of the average for their age. For non-white females, the poor die at a rate that is 186 percent of the average, while the rich die at a rate that is 44 percent of the average.

⁸ Income-differentiated mortality rates are also employed by Caldwell et al. (this volume) and Lee, McClellan, and Skinner (this volume).

Although we have the annual income of each individual in our sample for each year, we do not just use the corresponding annual income group's *O/E* ratio from Table 2 for that person in that year. One problem with doing so is that annual income levels would have to be adjusted for inflation and growth to match up with the 1980 levels in Table 2. A second problem is that an individual with a steeply hump-shaped earnings profile (as in Figure 1) would have a probability of dying that fell dramatically during high-annual-income years and then rose again during low-annual-income years. We do not believe that the same individual's probability of death changes that rapidly with annual income, jumping over other individuals in the same age cohort whose annual incomes are not so volatile. Instead, the probability of dying is more likely affected by the individual's *lifetime* income group. To solve both of these problems, our procedure described in section A.3 is based on the relative ranking of each individual's lifetime income. A person in a particular percentile of the lifetime income distribution gets the *O/E* ratio of a person in the same percentile of the annual income distribution.⁹

5. SOCIAL SECURITY TAXES PAID

This section describes how our calculation of social security tax for each person in each year follows the provisions of the Social Security Administration. This tax is commonly referred to as the FICA tax (Federal Insurance Contributions Act). It is collected on earned income and consists of three portions: Old Age and Survivors Insurance (OASI), Disability Insurance (DI), and Hospitalization Insurance (HI), also known as Medicare. The proceeds from these taxes are deposited into three separate trust funds, and benefits are paid from the appropriate fund. The program has become almost universal—95 percent of all employment in the U.S. is covered.¹⁰

The tax is deducted from employees' pay at a rate of 7.65 percent of wages, but employers match those deductions for a total tax of 15.3 percent. Self-employed individuals pay the entire 15.3 percent tax annually with their income tax returns. Both the employee and the employer

⁹ Thus, even if two retirees have the same low annual income, the one with higher lifetime income is assumed to have a lower mortality probability.

¹⁰ Coverage may be excluded for: federal civilian workers hired before 1984 who have not elected to be covered; railroad workers who are covered under a similar but separate program; certain employees of state and local government, covered by their state's retirement programs; household workers and farm workers with certain low annual incomes; persons with income from self-employment of less than \$400 annually; and those who work in the underground, cash, or barter economy who may illegally escape the tax.

share of the tax are collected on wages up to an annual maximum amount of taxable earnings—the social security wage cap (\$68,400 for 1998). This cap is adjusted automatically each year with the average earnings level of individuals covered by the system, thereby taking account of both real wage growth and inflation.

Since an objective of our research is to measure each worker's return on social security taxes, the question arises: how much of the total tax does the worker bear? Using only the statutory incidence (the worker's half) would result in much higher returns than using the combined employer and employee portions. Hamermesh and Rees (1993, p. 212) review empirical work on payroll tax incidence and conclude that the worker bears most of the employer's share of the tax through reduced wages. We therefore base our estimates on the combined employer and employee tax.¹¹

Our focus is the retirement portion of the social security system, not the disability insurance or hospital insurance. Of the total 15.3 percent employer and employee tax, 10.6 percent is for OASI, 1.8 percent is for DI, and 2.7 percent is for HI.¹² The OASI portion of the tax is paid directly to the OASI Trust Fund, which is used to pay all retirement benefits. We therefore ignore the DI and HI portions of the tax, as well as benefits paid from the DI and HI Trust Funds.

Table 3 shows the combined OASI tax rate for selected years since 1940, ending with the 10.6-percent rate for 2000 and beyond (as used in this study). The next two columns show the wage cap and the maximum possible tax.

In our study, we calculate the present value at age 22 of mortality-adjusted social security taxes and benefits through age 99. The probability P_{ij} of the individual being alive at age j is conditional on being alive at age 22, and it is computed from the constructed tables (for each age–race–sex–income cell) by

$$P_{ij} = \frac{\text{number in cell } i \text{ alive at age } j}{\text{number in cell } i \text{ alive at age 22}}$$

¹¹ Paris and Lillard (1996) point out that because the employer's portion of the payroll tax is deductible against the employer's income tax, the net cost to the employer is lower than the full amount of the payroll tax paid. Like Paris and Lillard, however, and for comparability with other studies, we treat the entire amount of the payroll tax as the employee's cost of social security coverage. In effect, we look at the social security system only, without any income tax. The combined incidence is not equal to the sum of the parts, but we cannot say whether the income tax affects the incidence of social security, or social security affects the incidence of the income tax.

¹² These allocation percentages are for the year 2000 and beyond. Congress "temporarily" increased the portion going to DI for the years 1994–1996, followed by a reduction for 1997–1999. The 1998 allocation is: OASI 10.7, DI 1.7, and HI 2.9 percent.

TABLE 3
Old Age and Survivors' Insurance (OASI) Tax Rates and Wage Caps

Year	OASI tax rate (%)	Taxable wage ceiling (\$)	Maximum tax (\$)
1940	2.00	3,000	60
1950	3.00	3,000	90
1960	5.50	4,800	264
1970	7.30	7,800	569
1980	9.04	25,900	2,341
1990	11.20	51,300	5,746
1995	10.52	62,700	6,596
2000	10.60	75,500	8,003

The tax rate is combined employer and employee portions. Wage ceiling in 2000 assumes constant inflation of 3% and wage growth of 1%.

Our sample from the PSID includes observed and constructed earnings for each individual from age 22 through the age of retirement. To obtain steady-state taxes and benefits under current law, however, we look at a hypothetical future cohort with a birth year of 1990. We therefore take N_{oij} , the *observed* nominal earnings of individual i in year j , and we convert it to the corresponding future-cohort individual's nominal earnings, N_{fij} , using the ratio of projected average earnings in the future year (AE_{fj}) to observed average earnings in the PSID sample year (AE_{oj}):

$$N_{fij} = N_{oij} \frac{AE_{fj}}{AE_{oj}}$$

Since 1951, the Social Security Administration has computed Average Earnings, the average annual earnings of all workers covered under the Act. We project this Average Earnings into the future, using assumptions about future real wage growth and inflation.¹³ Next, to compute SST_{ij} , the social security tax of person i in year j , we take

$$SST_{ij} = T \times \min(N_{ij}, CAP_j)$$

where T is the combined OASI tax rate (which is constant with unchanged law), and CAP_j is the maximum nominal earnings subject to the

¹³ We use actual inflation and growth to scale observed PSID years up to 1995. Since amounts in simulated future years are indexed, the subsequent inflation and growth rates are set to zero.

OASI tax (which increases with inflation). Then the amount that the individual expects (at age 22) to pay in year j is

$$E_{22}(\text{SST}_{ij}) = \text{SST}_{ij} \times P_{ij}.$$

That is, the future tax is only paid with the probability P_{ij} that person i is alive at age j . These amounts are used below, either to calculate an internal rate of return, or to calculate a present value using a particular discount rate.

6. SOCIAL SECURITY BENEFITS

Under provisions of the Social Security Act, benefits are calculated from a progressive formula based the individual's average indexed monthly earnings (AIME). Our calculations follow the Social Security Administration's computation of AIME upon the individual's retirement. In particular, earnings prior to age 60 are indexed to average wages in the year the individual attains age 60. The method of indexing is to multiply the nominal earnings in year j by the ratio of Average Earnings in the year age 60 was attained to Average Earnings in year j . Earnings after age 60 are not indexed. A person who worked from age 22 through age 64 (retiring on his or her 65th birthday) would have a total of 43 years of earnings. Under the Act, only the highest 35 years are considered, so the lowest eight years of earnings will be dropped. AIME is the simple average of the indexed earnings in those 35 highest-earnings years.¹⁴

Next, the primary insurance amount (PIA) is calculated as 90 percent of AIME up to the first bend point, plus 32 percent of AIME in excess of bend point 1 but less than bend point 2, plus 15 percent of AIME in excess of bend point 2. In 1995, bend point 1 was \$426 and bend point 2 was \$2,567. If AIME were \$3,200, for example, the PIA would be

$$\begin{aligned} \text{PIA} &= 0.90 \times (426) + 0.32 \times (2,567 - 426) + 0.15 \times (3,200 - 2,567) \\ &= \$1,163.47. \end{aligned}$$

Like the cap on earnings, the bend points are adjusted annually by the proportional increase in Average Earnings. We calculate this PIA for

¹⁴ The language of the Act specifies dropping the lowest *five* years of earnings through age 61. Then, if the worker has years of earnings after age 61 that are higher than some of the undropped years of earnings before age 62, the higher post-61 earnings will replace the lower pre-62 earnings. The net effect for a worker retiring at age 65 is to drop the lowest eight pre-65 years.

each worker in the sample, which then becomes the basis for all social security benefit calculations.

A retiree is entitled to a benefit equal to the PIA upon normal retirement at age 65. Legislation already enacted will increase the retirement age by two months each year beginning in 2000, so that by 2005 the normal retirement age will be 66. Another two-month-per-year increase will begin in 2017, resulting in a normal retirement age of 67 after the year 2021. A worker may still choose to retire as early as age 62, with reduced benefits.¹⁵ In contrast, if a worker elects to *delay* receipt of benefits to an age as late as 70, the eventual benefits are permanently increased by 5% per year of delay.¹⁶ Our calculations below ignore these provisions for early or late retirement, as we assume workers (and their spouses) always choose the normal retirement age.¹⁷

In addition to retirement benefits for covered workers, the OASI Trust Fund provides certain benefits to the spouse and other dependents of retired or deceased workers. In the aggregate, the non-spousal survivor benefits are relatively minor.¹⁸ We take account of payments to a surviving spouse, but we ignore the non-spousal survivors benefits and therefore slightly understate the rates of return.

The spouse of a retired worker can receive the benefit based on his or her own earnings, or one-half of the PIA of the retired worker (designated as the "spousal benefit"), whichever is greater. The age at which a spouse may receive the full spousal benefit is the same normal retirement age as for a retired worker. The spouse may elect to receive the benefit as early as age 60, provided that the working spouse has retired.¹⁹ No premium is provided for delaying the spousal benefit, but we

¹⁵ This early-retirement penalty is a permanent reduction in the PIA of $\frac{5}{9}$ percent for each early month (6.67 percent for each early year). For example, a worker retiring at age 62 when the normal retirement age is 65 would receive a benefit for the rest of his or her life that is reduced by 20 percent.

¹⁶ Under the Act, the 5-percent-per-year premium is scheduled to increase gradually to 8 percent per year for workers reaching age 65 in 2008. The increase applies only for delayed retirement up to age 70. Beyond age 70, no further incentive is provided to delay benefits.

¹⁷ This assumption does not affect progressivity unless the chosen date of retirement differs by income. If low-income individuals tend to die earlier, then they might optimally retire earlier, so the availability of this option might be progressive.

¹⁸ In 1996, a total of \$302.9 billion in benefits were paid from the OASI trust fund. Of that total, \$288.1 billion was paid as retirement benefits to retired workers or their spouses, and only \$14.8 billion (4.9 percent) was paid for other survivor and miscellaneous benefits (*Annual Statistical Supplement—1997*, Table 4A.5 in U.S. Social Security Administration, 1998).

¹⁹ However, the penalty for receiving the spousal benefit is somewhat higher than the worker's penalty: $\frac{25}{36}$ percent for each of the first 36 months (8.33 percent per year) before

always assume normal retirement age anyway. Then, once spousal benefits have begun, cost of living adjustments for the spousal benefit are handled in the same manner as for the worker's benefit.

The spouse of a *deceased* worker can receive the benefit based on his or her own earnings, or 100% of the benefit to which the retired worker was entitled, whichever is greater. The benefit based on the deceased worker's benefit is called the *survivor benefit*.

Our calculations of these amounts are detailed in section A.4. Our PSID sample provides a complete history of observed and constructed earnings for each individual in the sample for ages 22 through 66. They then retire when they turn 67, the normal retirement age under current law for that future cohort of individuals. We use those earnings to compute indexed earnings, the AIME, the PIA, the spousal benefit (SpBen), and the survivor benefit for the surviving spouse (SurvBen) in exact accordance with provisions of the Act.

7. RESULTS FOR THE CURRENT SOCIAL SECURITY SYSTEM

We use procedures described above to calculate the mortality-adjusted tax and benefit in each year for each individual in each of our lifetime income quintiles. Then we use three different measures to gauge the distributional effect of the social security system (and later, proposed reforms). First, for each simulation, we compute the present value, at age 22, of the benefits to be received minus the taxes paid. We then sum over the individuals in each lifetime income quintile. This measure indicates the absolute size of the social security transfers between income groups. The discount rate should reflect a real rate of return that would be available to participants in the system, and that would provide for the same certainty as does the Social Security System. The Trustees of the Social Security System currently use a rate of 2.8 percent for their long-term estimate of real returns.²⁰ Ibbotson Associates (1998) reports on historic rates of return for various portfolio investments. For the period 1935 to 1997, the average inflation rate was 4.0 percent, and the nominal

normal retirement age, plus $\frac{5}{12}$ percent (5 percent per year) for each of up to 24 additional months before normal retirement age. A 62 year-old wife of a retired worker would be entitled to a spousal benefit equal to 75 percent of the normal 50 percent of the retired worker's PIA.

²⁰ In arriving at that rate, they forecast inflation at a long-term rate of 3.5 percent, and a nominal interest rate of 6.3 percent on the special-issue U.S. Treasury obligations that are purchased by the OASI trust fund. Whether to use a before-tax or an after-tax discount rate depends on one's assumption about what alternative retirement investments are available.

return on intermediate-term U.S. Treasury obligations was 5.4 percent, so the real rate of return was 1.4 percent.²¹

For one choice of discount rate we use 2 percent, which lies between the forecast rate earned by the OASI trust fund on its investments (2.8 percent) and the historical average of real returns on government bonds reported by Ibbotson (1.4 percent).²² To test the sensitivity of results, we also use a discount rate of 4 percent. As shown below, the choice of rate affects not only the absolute size of the present-value gains or loss for each group but also the pattern of progressivity.

Second, we also express the present value of net benefits as a percentage of the present value at age 22 of the lifetime endowment (discounted at the same rate). This scaling reveals the relative progressivity of the system. If the same absolute net benefit is provided to all individuals, that benefit will constitute a higher percentage of lifetime income in the lowest quintile. Such a system would typically be called progressive.

Third, we ignore the chosen discount rate, and we calculate the internal rate of return (IRR) that equates the present value of benefits with the present value of taxes. The IRR has become an almost universal tool for measurement of social security taxes and benefits, and we include it for comparability with other studies.

Our initial simulations use the enacted provisions of the Social Security Act, applied to a future cohort born in 1990. Results *without* income-differentiated mortality are presented in Table 4. For each quintile, this table shows the average undiscounted taxes paid and benefits received, as well as the present value of net benefits, and the internal rates of return from the streams of taxes and benefits. These results could be viewed as the distributional impact of an extreme reform—repeal of the whole social security system.

The reason for showing undiscounted taxes and benefits is to shed some light on the overall solvency of the social security system. Our model cannot project actual inflows and outflows, since we do not use demographic forecasts, but a conceptual point can be made about solvency in a world with unchanging demographics: with a constant num-

²¹ The nominal return on long-term Treasury obligations was actually lower, 5.3 percent, for a real return of 1.3 percent. Investing in U.S. Treasury Bills, which are the benchmark for risk-free investments, yielded a nominal return of 4.0 percent, and a real return of 0 percent.

²² Other studies of social security redistribution have used rates on either side of 2 percent. Myers and Schobel (1983) use 2 percent, Hurd and Shoven (1985) use 3 percent, Boskin, Kotlikoff, Puffert, and Shoven (1987) use 3 percent, Duggan, Gillingham, and Greenlees (1993) use 1.2 percent, Steuerle and Bakija (1994) use 2 percent, and Gramlich (1996) uses 2.3 percent. In contrast, Caldwell et al. (this volume) use 3, 5, or 7 percent.

TABLE 4
Lifetime Taxes and Benefits in the Base Case for the Average Person in Each Group
(without Income-Differentiated Mortality)

Quintile (1)	Taxes paid (undiscounted) (\$ thousands) (2)	Benefits received (undiscounted) (\$ thousands) (3)	Present value of net benefits (%) (4)		Internal rate of return (%) (5)	(Present value of net benefits)/ (present value of lifetime income) (%) (6)	
			2%	4%		2%	4%
			Discount	Discount		Discount	Discount
1	64.7	125.7	-1.3	-16.7	1.92	-0.17	-3.10
2	85.4	151.8	-5.8	-22.7	1.69	-0.53	-2.94
3	106.0	168.7	-13.8	-30.4	1.37	-1.05	-3.26
4	115.8	174.8	-17.8	-33.6	1.23	-1.13	-3.00
5	141.4	187.0	-30.1	-43.3	0.85	-1.33	-2.70
All	102.7	161.6	-13.8	-29.4	1.35	-0.98	-2.96

Quintile 1 is the lowest lifetime income group. Columns 2 and 3 are undiscounted lifetime taxes paid and benefits received (1995 \$). Column 4 is the present value at age 22 of benefits less taxes, discounted at 2% and 4%. Column 5 is the internal rate of return from the net benefits received.

ber of entering 22-year-olds in each of the sex-race-income cells in our model, the undiscounted sum of taxes paid for an individual (\$102,700) equals the sum of taxes paid by all ages alive at one time. Similarly the undiscounted sum of benefits (\$161,600) is the sum of benefits paid out to all ages alive at one time. On this basis, the pay-as-you-go social security system loses the difference (\$58,900 per 22-year-old) each year.²³

Table 4 shows the progressivity of the current social security system. Individuals in the lowest lifetime income quintile pay lifetime taxes of \$64,700 and can expect to receive lifetime benefits of \$125,700. Those benefits come later in life, however, so the present value at age 22 is a net loss of \$1,300 (using the 2-percent discount rate). This net loss is 0.17 percent of their discounted lifetime endowment. In contrast, the highest income group pays taxes of \$141,400 and receives benefits of \$187,000, but discounting net benefits at 2 percent to age 22 results in a present-

²³ If we multiply this \$58,900 figure by the number of 22-year-olds alive in 1994 (about 3.7 million), we get a total loss of about \$220 billion per year. This figure lies between the "low" and the "high" deficit projected by the Board of Trustees of the Social Security System (U.S. Social Security Administration, 1998). When converted into 1995 dollars, their "intermediate" projected deficit for 2075 is \$480 billion, but that includes DI and pertains to a larger population.

value loss of \$30,100, which represents 1.33 percent of the present value of their lifetime income.

The internal IRR is 1.92 percent for the lowest income group and 0.85 percent for the highest income group.²⁴ Viewing these results, one might say that the social security system is highly progressive. Unfortunately, however, the IRR can be very sensitive to small changes in parameter values. Moreover, the tax-incidence literature usually defines progressivity by an average tax rate that rises with income. In our model, the lifetime average tax rate is the present value of net tax divided by the present value of income. This ratio increases from 0.17 percent for the lowest income group to 1.33 percent for the highest income group. By the usual definition, then, the social security system is progressive.

When the discount rate is raised to 4 percent, the taxes in working years become even larger relative to the benefits received in retirement years. The present-value net loss rises from \$1,300 to \$16,700 for the lowest income quintile and from \$30,100 to \$43,300 for the highest quintile. These losses are all now close to 3 percent of the present value of lifetime endowment.²⁵

What differences do we anticipate with the introduction of income-differentiated mortality? Since individuals with low lifetime incomes have shorter than average life spans, they will receive benefits for a shorter period. Thus we expect less progressivity. Results with income-differentiated mortality are shown in Table 5. For a rough measure of progressivity, consider the difference between the two groups' net tax rates. In the non-differentiated results of Table 4, the net tax rate was 0.17 percent for the lowest quintile and 1.33 percent for the highest quintile, for a difference of 1.16 percent. In the income-differentiated results of Table 5, these numbers are 0.60 and 1.01 percent, for a difference of only 0.41 percent.²⁶ Since the measure of progressivity falls from 1.16 to 0.41, we conclude that the consideration of income-differentiated mortality has reduced the progressivity of the social security system by more than half.

²⁴ Caldwell et al. (this volume) also find that the IRR falls from their lowest to their highest lifetime-labor-earnings group.

²⁵ These results are consistent with those of Caldwell et al. (this volume). Because we divide the net tax by the present value of potential earnings for 4,000 hours per year, our net tax rates are lower than those they get when they divide the net tax by the present value of actual earnings.

²⁶ At the same 2-percent discount rate but without income-differentiated mortality, in Table 4, the ratio of net tax to lifetime income rises monotonically with income. When income-differentiated mortality is introduced in Table 5, the ratio is hump-shaped. It rises from the lowest to the middle group and then falls for the highest two groups.

TABLE 5
Lifetime Taxes and Benefits in the Base Case for the Average Person in Each Group
(with Income-Differentiated Mortality)

Quintile (1)	Taxes paid (undiscounted) (\$ thousands) (2)	Benefits received (undiscounted) (\$ thousands) (3)	Present value of net benefits (%) (4)		Internal rate of return (%) (5)	(Present value of net benefits)/ (present value of lifetime income) (%) (6)	
			2%	4%		2%	4%
			Discount	Discount		Discount	Discount
1	63.2	113.9	-4.5	-17.6	1.69	-0.60	-3.31
2	84.8	144.8	-7.8	-23.3	1.58	-0.72	-3.03
3	106.5	169.5	-13.8	-30.4	1.38	-1.05	-3.26
4	117.3	186.0	-15.0	-32.8	1.38	-0.94	-2.91
5	143.9	213.1	-23.1	-41.3	1.17	-1.01	-2.55
All	103.1	165.5	-12.8	-29.1	1.40	-0.91	-2.92

Quintile 1 is the lowest lifetime income group. Columns 2 and 3 are undiscounted lifetime taxes paid and benefits received (1995 \$). Column 4 is the present value at age 22 of benefits less taxes, discounted at 2% and 4%. Column 5 is the internal rate of return from the net benefits received.

An increase in the discount rate cuts progressivity by more and can even make the overall social security system regressive. In the last column of Table 5, with 4 percent discount rate, the net tax rate on the lowest-income quintile (3.31 percent) is higher than on the top-earning quintile (2.55 percent). In general, a higher discount rate reduces the present value of progressive benefits received during later retirement years by more than it reduces the present value of regressive taxes paid during earlier working years.

8. PROPOSED REFORMS TO THE SOCIAL SECURITY SYSTEM

A large number of recent articles on social security reform have dealt with privatization of the system or other large-scale overhauls (e.g. Kotlikoff, Smetters, and Walliser, 1998). If complete privatization were to provide actuarially-fair returns, with no redistributions between individuals, then the effects of complete privatization in our model are exactly the reverse of *having* the current social security system (in Tables 4 and 5).

However, political considerations may preclude radical reforms. Be-

cause of the currently-projected deficits in the long run, realistic reforms might just require one or more changes to raise taxes or reduce benefits within the context of the current system. We therefore consider piecemeal reforms like increasing the retirement age, changing the manner in which the benefit is computed, increasing the payroll tax, or decreasing the overall level of benefits. All of these reforms were considered by the 1994–1996 Advisory Council on Social Security (1997).²⁷

8.1 Elimination of the Drop-Years Provision

As discussed above, a worker's AIME is computed by adjusting each year's earnings for inflation and real wage growth, and then dropping the lowest five years of pre-age-62 indexed earnings. In addition, if the individual works beyond that age, non-indexed earnings for those additional years may be substituted for any lower year's earnings. Thus, for a worker retiring at age 65, a total of eight low-earnings years may be dropped. By 2020, when the normal retirement age is 67, a total of ten years may be dropped.

The effect of the "drop" provision is to increase AIME and therefore to increase benefits most for individuals who have high variability in their lifetime earnings pattern. For example, suppose one individual had level earnings for all working years, while another had the same earnings each year except for ten years with no earnings. The first individual pays substantially more social security tax but receives the same benefit as the second individual who took ten years off.

Steuerle and Bakija (1994) assume that higher-earning individuals have higher earnings variation than do lower-earnings individuals, and they point out that this drop-years provision is not likely to be progressive.²⁸ Williams (1998) has similar findings. The model in this paper can be used to evaluate this point about the progressivity of the drop-years provision.

Our first simulated reform is to delete the drop-years provision, which

²⁷ The Advisory Council issued a list of consensus recommendations and three non-consensus sets of other recommendations. Consensus recommendations included two of our proposed reforms—increase the period over which wages are averaged in the benefit formula and increase the normal retirement age beyond 67. Two of the three non-consensus groups proposed various types of benefit reductions (another of our reforms), while one of the groups proposed increasing the payroll tax (our final reform).

²⁸ "The [dropout-years provision] is hardly progressive. Among its beneficiaries are those who delay entrance into the full time labor force to attend college and graduate or professional school. . . . Dropout years especially tend to discriminate against those lower-wage workers who enter the labor force at a younger age and stay in the labor force throughout most of their adult lives" (Steuerle and Bakija, 1994, p.185).

TABLE 6
Deleting the Drop-Years Provision: Changes from the Base-Case Simulation

Quintile (1)	Without income-differentiated mortality			With income-differentiated mortality		
	Present value of change in net benefits (\$ thousands) (2)	(2) divided by lifetime income (%) (3)	Change in internal rate of return (%) (4)	Present value of change in net benefits (\$ thousands) (5)	(5) divided by lifetime income (%) (6)	Change in internal rate of return (%) (7)
1	-5.7	-0.75	-0.41	-5.1	-0.68	-0.40
2	-5.5	-0.51	-0.32	-5.2	-0.48	-0.33
3	-4.7	-0.36	-0.25	-4.6	-0.35	-0.25
4	-4.9	-0.31	-0.25	-5.1	-0.32	-0.26
5	-5.7	-0.25	-0.29	-6.6	-0.29	-0.28
All	-5.2	-0.37	-0.30	-5.4	-0.38	-0.29

Quintile 1 is the lowest lifetime income group. Columns 2 and 5 are the changes in the present value at age 22 of benefits minus taxes paid, discounted at 2%. Columns 3 and 6 equal columns 2 and 5 divided by the present value at age 22 of lifetime income discounted at 2%. Columns 4 and 7 are the changes in the internal rates of return of social security taxes and benefits.

means that all earnings for ages 22 through 66 are included in the AIME computation. This reform reduces every individual's net benefit and internal rate of return. Results for our five lifetime income groups are shown in Table 6. Without income-differentiated mortality, in the left-hand side of Table 6, the absolute decline in net benefits is fairly flat across these groups, falling by \$5,700 in both the lowest income group and the highest income group. In relative terms, however, the effect is fairly regressive. The same \$5,700 decline in net benefits represents 0.75 percent of lifetime income for the lowest quintile and only 0.25 percent of lifetime income for the highest quintile (a regressive spread of 0.50 percent). We also examine results by marital status and gender, finding the same pattern throughout: the elimination of the drop-years provision is somewhat regressive (across income groups) for all demographic categories.

With income-differentiated mortality, in the right-hand side of Table 6, this reform appears slightly less regressive. The lowest quintile has a decrease in net benefits equal to 0.68 percent of lifetime income, while the highest quintile has a decrease of 0.29 percent (a regressive spread of only 0.39 percent).

To explain the difference between anticipated and actual results,

TABLE 7
Effect of Reduction in the Number of Drop Years from 10 to 0

Quintile (1)	Decline in AIME (%) (2)	Decline in lifetime benefits (%) (3)
1	-19.8	-13.3
2	-18.5	-10.5
3	-17.1	-8.1
4	-16.9	-8.1
5	-17.7	-9.0
All	-17.8	-9.6

Quintile 1 is the lowest lifetime income group. Column 2 is the percentage decline in AIME, and column 3 is the percentage decline in undiscounted lifetime benefits that result from eliminating the drop-years provision.

Table 7 shows the change in AIME and the change in lifetime benefits for each quintile. Eliminating the drop-years provision causes a reduction in the AIME calculation for low-income individuals of 19.8 percent and a reduction for the highest income group of 17.7 percent. Thus, variability of earnings is not materially different. Since the benefit formula is progressive, however, the reduction in AIME for high-income individuals generates a lower percentage decrease in their benefits than for lower quintiles. Thus, approximately equal proportional reductions in AIME result in a larger percentage cut in benefits for low-income workers than for high-income workers.

8.2 *Increase in Retirement Age*

The adoption of age 65 as the normal retirement age for social security can be traced to nineteenth-century Germany. Steuerle and Bakija (1994) report that when the first German universal retirement system was introduced, Bismarck apparently chose an age beyond which few survived. Since the Social Security Act was passed in 1936, the average life expectancy has increased by almost four years, and the Social Security Administration expects the trend to continue into the next century. A male who was 65 in 1936 could expect to live to age 76.6 (11.6 years beyond age-65 retirement), whereas his counterpart in 2026 can expect to live to 81.6 (16.6 years of retirement).²⁹

²⁹ If medical technology tends to increase the quality of life as well as the length of life expectancy, then it might be natural to suppose that people can work to older ages as well as live to older ages.

Already-enacted amendments to the Social Security Act will gradually increase the normal retirement age to 67 by the year 2020.³⁰ We now analyze a reform that would further increase the retirement age beyond the scheduled increase to age 67—thereby increasing the duration of tax collections and decreasing the duration of benefits.

How much to increase the retirement age? We want to analyze and compare the distributional effects of reforms that have the same overall impact on social security. Each reform may reduce benefits or raise taxes, or both, so we standardize all reforms to have the same net impact on the social security shortfall as the elimination of the drop-years provision. As shown above, the annual net shortfall is the sum of undiscounted benefits (paid out to all ages alive at one time) minus undiscounted taxes (received from all ages alive at one time). In Table 4, the current system's shortfall is \$58,900 per 22-year-old. As it turns out, the elimination of the drop-years provision reduces the average undiscounted net benefits from \$58,900 to \$43,400. We then use that net revenue gain (\$15,500) as our target for the other reforms.³¹ To achieve this target through a change only to the retirement age, we increase the retirement age from 67 to 68.4 years.³²

We further assume that individuals will continue to work until normal retirement age, even when the normal retirement age is increased. Since the number of years in the work force is increased and the expected number of years of drawing benefits is reduced, this reform results in higher taxes *and* lower benefits—a reduction in net benefits and in the internal rates of return for all workers.

Table 8 shows that the present value of net benefits falls more for high-income groups. The 1.4 years of additional work and delay of benefits hits the higher income groups harder in absolute terms. When we divide the decline in net benefits by the income of each group, however, the distributional results are shown to be regressive. To explain this

³⁰ Workers will still be able to retire as early as age 62, but with a larger penalty. While a current age-62 retiree receives 80 percent of the full benefit, the eventual age-62 retiree receives only 70 percent.

³¹ Again, our measure of revenue is somewhat stylized, since we do not have detailed demographic projections. It is the real revenue that would be collected in any one year of the future steady-state growth path with *unchanging* demographics, summing over all income groups and ages alive in that year.

³² The same target can be achieved by a 15.1-percent increase in the tax rate or by a 9.6-percent decrease in all benefits. With income-differentiated mortality, the current system's shortfall was \$62,300 (per 22-year-old). Eliminating the drop-years provision reduced that shortfall to \$46,600, so the net revenue target in that case is \$15,700. This target is also met by an increase in the retirement age from 67 to 68.4, by a 15.3-percent increase in the tax rate, or by a 9.5-percent decrease in all benefits.

TABLE 8
*Increasing the Normal Retirement Age from 67 to 68.4:
 Changes from the Base-Case Simulation*

Quintile (1)	Without income-differentiated mortality			With income-differentiated mortality		
	Present value of change in net benefits (\$ thousands) (2)	(2) divided by lifetime income (%) (3)	Change in internal rate of return (%) (4)	Present value of change in net benefits (\$ thousands) (5)	(5) divided by lifetime income (%) (6)	Change in internal rate of return (%) (7)
1	-4.9	-0.64	-0.34	-4.5	-0.60	-0.33
2	-5.8	-0.54	-0.33	-5.7	-0.52	-0.34
3	-6.5	-0.49	-0.32	-6.5	-0.50	-0.33
4	-6.8	-0.43	-0.33	-6.9	-0.44	-0.33
5	-7.3	-0.32	-0.33	-8.0	-0.35	-0.31
All	-6.2	-0.45	-0.33	-6.4	-0.45	-0.33

Quintile 1 is the lowest lifetime income group. Columns 2 and 5 are the changes in the present value at age 22 of benefits minus taxes paid, discounted at 2%. Columns 3 and 6 equal columns 2 and 5 divided by the present value at age 22 of lifetime income discounted at 2%. Columns 4 and 7 are the changes in the internal rates of return of social security taxes and benefits.

overall regressive result, note that individuals pay the regressive tax for a longer period and receive the progressive benefit formula for a shorter period.

Comparing the left-hand and right-hand sides of the table shows that consideration of income-differentiated mortality reduces somewhat the regressive nature of this reform.

8.3 Increase in the Social Security Tax Rate

Prior to 1990, the social security tax rate was increased frequently—a total of 20 times from the inception of the system. Some of these increases are shown in Table 3. Amendments to the Act in 1983 provided for a constant total tax rate of 15.3 percent from 1990 onward, and Congress appears very reluctant to increase the tax beyond that level. Nonetheless, the target reduction in the social security shortfall could be achieved by a simple increase in the rate of tax.

To raise net revenue by the same amount as the first two reforms, the OASI portion of the payroll tax must be increased by 15.1 percent (from 10.6 to 12.2 percent). This reform leaves benefits unchanged.

As shown in Table 9, this reform imposes higher additional tax on higher income groups. As a fraction of lifetime income, however, the

TABLE 9
Increasing the Tax Rate by 15.3%: Changes from the Base-Case Simulation

Quintile (1)	Without income-differentiated mortality			With income-differentiated mortality		
	Present value of change in net benefits (\$ thousands) (2)	(2) divided by lifetime income (%) (3)	Change in internal rate of return (%) (4)	Present value of change in net benefits (\$ thousands) (5)	(5) divided by lifetime income (%) (6)	Change in internal rate of return (%) (7)
1	-6.6	-0.87	-0.41	-6.6	-0.88	-0.40
2	-8.7	-0.80	-0.41	-8.7	-0.80	-0.42
3	-10.8	-0.82	-0.41	-10.9	-0.83	-0.42
4	-11.7	-0.74	-0.42	-11.9	-0.75	-0.43
5	-14.1	-0.62	-0.43	-14.6	-0.63	-0.42
All	-10.3	-0.73	-0.42	-10.6	-0.75	-0.42

Quintile 1 is the lowest lifetime income group. Columns 2 and 5 are the changes in the present values at age 22 of benefits minus taxes paid, discounted at 2%. Columns 3 and 6 equal columns 2 and 5 divided by the present value at age 22 of the lifetime income discounted at 2%. Columns 4 and 7 are the changes in the internal rates of return of social security taxes and benefits.

change is somewhat regressive. Note that social security always combines a regressive payroll tax and a progressive benefit formula. The net effect could go either way. In this reform, all individuals pay more of the regressive tax with no change in benefits, so the net effect is regressive.

With income-differentiated mortality, the right-hand side of Table 9 shows that this reform is about equally regressive. Mortality assumptions matter less for this reform because taxes are paid earlier in life when mortality rates are smaller.

8.4 *Decrease in the Overall Benefit Level*

The social security benefit formula is progressive: 90 percent of AIME up to bend point 1, 32 percent of AIME between bend points 1 and 2, and 15 percent of AIME, if any, beyond bend point 2. These rules for social security benefits were amended frequently prior to the 1977 amendments. Since 1977, the only increases in benefits have resulted from automatic inflation and real-wage-growth increases in the bend points of the benefit formula. Nevertheless, one approach to cutting the shortfall in the social security system is to decrease the overall level of benefits.

If every annual benefit were reduced by 9.6 percent, then the annual revenue shortfall would be reduced by the same amount as in the first

TABLE 10
Decreasing the Benefit Level by 9.5%: Changes from the Base-Case Simulation

Quintile (1)	Without income-differentiated mortality			With income-differentiated mortality		
	Present value of change in net benefits (\$ thousands) (2)	(2) divided by lifetime income (%) (3)	Change in internal rate of return (%) (4)	Present value of change in net benefits (\$ thousands) (5)	(5) divided by lifetime income (%) (6)	Change in internal rate of return (%) (7)
1	-4.1	-0.54	-0.29	-3.6	-0.48	-0.28
2	-5.0	-0.46	-0.29	-4.7	-0.43	-0.30
3	-5.5	-0.42	-0.29	-5.4	-0.41	-0.30
4	-5.7	-0.36	-0.30	-6.0	-0.38	-0.30
5	-6.1	-0.27	-0.31	-6.9	-0.30	-0.29
All	-5.2	-0.37	-0.30	-5.4	-0.38	-0.29

Quintile 1 is the lowest lifetime income group. Columns 2 and 5 are the changes in the present values at age 22 of benefits minus taxes paid, discounted at 2%. Columns 3 and 6 equal columns 2 and 5 divided by the present value at age 22 of the lifetime income discounted at 2%. Columns 4 and 7 are the changes in the internal rates of return of social security taxes and benefits.

three reforms. Table 10 shows that this reform would provide the low-income quintile with the smallest cut in the dollar amount of net benefits received. In relation to income, however, the change is regressive. An individual in the lowest lifetime income group gives up net benefits equal to 0.54 percent of lifetime income, while one in the highest group gives up only 0.27 percent of lifetime income. The regressive nature of this change is inherent, since the regressive tax is unchanged while the progressive benefit formula is reduced. With income-differentiated mortality, however, the change is not *as* regressive.

9. CONCLUSION

Social security redistributes not only from young to old, but from rich to poor. The amount of that latter redistribution is the subject of this paper. We look at a large sample of individuals from the PSID, estimate their lifetime potential earnings, and categorize them into quintiles. We also use observed actual earnings in the sample years and construct earnings in other years to obtain an entire earnings history for each individual.

These histories are used to calculate social security taxes in each working year and benefits in each retired year.

Standard mortality tables stop at age 85, and they are not differentiated by income. In this paper, we extend mortality tables to 99 years of age, and we use evidence on mortality differences by annual income groups to develop tables that differ among our lifetime income groups. Without income-differentiated mortality, the social security system is fairly progressive across our five lifetime income categories. With our income-differentiated mortality tables, however, a major portion of the progressivity of the social security system disappears. Remaining progressivity can be reduced or reversed by an increase in the discount rate.

Finally, we analyze four reforms that would raise the same amount of revenue. Since social security taxes are regressive and benefits are progressive, however, any across-the-board increase in tax or decrease in benefits is a regressive change. The consideration of income-differentiated mortality somewhat reduces the regressivity of the reforms.

APPENDIX: DATA AND METHODOLOGY

This appendix is divided into four parts, describing the selection of the sample from the PSID, the estimation of earnings profiles, the derivation of income-differentiated mortality, and the calculation of social security benefits.

A.1 Data

We select our sample from the PSID based on three criteria. First, our sample members are not taken from the low-income subsample of the PSID. While the data contain weights so that the low-income sample can be merged with the representative sample, we felt that the representative sample provided sufficient data for our purposes. Second, we require that sample members remain in the sample for the entire period. Survey respondents may have died, or simply decided that the survey was no longer worth their time. Including those who dropped out of the sample was judged not to be worth the possible distortion in the data and additional computational work required to track these individuals. Third, we require that sample members do not change marital status during the sample period. It is difficult to incorporate changes in marital status in an analysis of lifetime incomes. We thus decided to include only those individuals whose status did not change, despite the biased sample selection this implies.

As our analysis is intended to reflect a steady state, we abstract from real economic growth that occurred during our sample period. We want

to isolate life-cycle movements in wages so that our wage profiles are not specific to one generation during a particular time frame. Adjusting for economic growth and inflation yields lifetime wage profiles that can be used to analyze the distributional impact of social security in a more general, structural sense. We therefore adjust the nominal wage rate using the Social Security Administration's Average Wage Index, which reflects growth in average nominal wages over the sample period. Using this index to deflate wages thus removes the effects of both inflation and real growth in wages.

We want to estimate a wage regression for the working wives as we do for the heads, but we question the idea of pooling the positive observations of the wives who work consistently throughout the sample with those who work only occasionally. We found that a woman would have to work at least 750 hours a year throughout her working life, an amount slightly less than half time, to have her own social security benefits be greater than the spousal benefits she could receive based on her husband's earnings (assuming she earns the same wage as her husband). Thus, we divide the working wives into two groups based on whether or not they averaged at least 750 hours of work per year throughout the sample. We ran our log wage regressions separately for the two groups, and then ran another one pooling the two groups, in order to perform an F test. The results suggest that these two groups should indeed be analyzed separately. These regressions are described in the text.

A.2 The Estimation of Earnings Profiles

For each of our lifetime income quintiles, we estimate separate earnings regressions for heads, habitual working wives, and part-time working wives. Our dependent variable is actual annual earnings. As in the first stage, we deflate earnings by the Social Security Administration's Average Wage Index to adjust for both inflation and real economic growth. Since earnings represent a continuous variable truncated at zero, we use a tobit framework for estimation. Here we assume that earnings is the product of optimal hours of work and a wage rate that is exogenous to the individual. Optimal hours of work can be positive or negative, so optimal earnings can be described as a latent variable y^* :

$$y_i^* = X_i\beta + \epsilon_i \quad ,$$

where X is a vector of personal characteristics that determine the individual's wage and desired hours of work. We assume that observations of zero hours worked imply that desired hours of work are less than or

equal to zero. Actual earnings, y , are observed only if y^* is greater than zero. If y^* is less than or equal to zero, then actual earnings are zero:

$$y_i = \begin{cases} y_i^* & \text{if } y_i^* > 0, \\ 0 & \text{if } y_i^* \leq 0. \end{cases}$$

In the first stage described above, in which we divide people into lifetime income quintiles, our dependent variable was log wages. Thus we use generalized least-squares estimation with individual fixed effects. In this second stage, the tobit model is nonlinear. We judged that the additional programming effort to include fixed effects in our tobit estimation was not worthwhile, given that such estimation also implies inconsistent parameter estimates (Heckman and MaCurdy 1980). By excluding fixed effects in this stage, we are able to include race, gender, and education variables in the earnings regressions without interacting them with age.

To simulate out-of-sample observations, we multiply the independent variables of each individual by the appropriate coefficients from their group's earnings regression. In addition, we include a random component, which we obtain by using the estimated standard error of each group's regression to generate a normally-distributed random variable. This random component is intended to represent unforeseen circumstances that affect earnings. It also means that individuals with the same observed characteristics will not have the exact same earnings profile. Simulated earning observations are thus calculated as

$$\hat{y}_i = X_i \hat{\beta} + \hat{\epsilon}_i,$$

where $\hat{\beta}$ is the vector of estimated coefficients from our earnings regressions, and $\hat{\epsilon}$ is the random component obtained by using the standard error of the regression to generate a random variable. Using this procedure, both positive and zero observations are generated. We found that the number of zeros generated for each group is consistent with the number of zero observations observed for that group during the sample years.

A.3 Derivation of Extended, Income-Differentiated Mortality

To extend the mortality tables from age 85 through age 99, we make three assumptions. First, we assume that the probability of remaining alive beyond age 85 decreases annually by a constant amount (Faber and Wade, 1983). Second, we set to zero the probability of remaining alive after age 99. This age seems a reasonable cutoff point, since less than 0.7 percent of

all social security beneficiaries are older than 95 (*Annual Statistical Supplement*, 1995). Third, given these two conditions, we find the constant annual change in the probability each year for each sex-race group such that the resulting set of probabilities yields the same life expectancy at age 85 as in the *Vital Statistics*. The result was shown in Figure 2.

Table 7 in Rogot, Sorlie, Johnson, and Schmitt (1992) shows information on actual deaths in their sample for each annual income group, within each race-sex-age group. For example, consider white males, ages 25 to 34. For each range of income (e.g. \$10,000 to \$14,999 in 1980 dollars), their table shows the number of individuals in their sample ($N = 14,563$), the number of observed deaths during the sample period ($O = 115$), and the number of deaths that would be expected if all income groups had the same mortality rate ($E = 92.2$). They then divide to get the observed/expected ratio ($O/E = 1.25$). Actual deaths in that low-income group are 25 percent higher than what would be expected using tables not differentiated by income.

We know the annual income of every individual in our PSID sample, so we need to exclude the "unknown income" category from the table in Rogot, Sorlie, Johnson, and Schmitt (1992). If we simply ignored this category, the overall O/E ratio would not be 1.0 for all income groups together. For this reason, we recalculate the expected deaths based on the subset of their individuals for which income is known, and recalculate O/E ratios for each group. The average of these new O/E ratios is 1.0, as desired. We then apply the appropriate ratio to each cell. Results for 25-34-year olds are shown in Table 2.

Finally, since annual income is volatile, we do not want to apply these annual-income-differentiated O/E ratios to the annual income of each person each year. Instead, we base differential mortality on lifetime income, in three steps. First, after we compute the present value of lifetime income for each of the 1,786 in our PSID sample, we assign each individual a ranking compared to all individuals in our sample. For example, an individual whose lifetime income ranks 432 out of the 1,786 individuals is ranked in the 24th percentile. Second, for each of the annual income groups in Table 2, we likewise determine percentile rankings based on income (shown in the third column). Third, for each individual in our sample, we match the percentile of their lifetime income to the percentile for the same age-race-sex category in Table 2. For example, a white female aged 27 who has lifetime income at the 24th percentile would be matched to the \$10,000-14,999 annual-income group (which lies between the 18th and the 36th percentile). That individual would then be assigned that group's O/E ratio for white females (1.17). Finally, this ratio is used to scale the probability of death for that

individual's age, sex, and race in the *Vital Statistics* (which are not differentiated by income).

A remaining problem, however, is related to causality: our procedure essentially uses the individual's income as a determinant of death, even though the annual income levels in Table 2 may be determined in part by illness immediately preceding death. This problem is somewhat mitigated by the fact that the CPS data used by Rogot, Sorlie, Johnson, and Schmitt (1992) are based on total combined family income, rather than just the decedent's income.

A.4 Calculation of Social Security Benefits

Every variable in this appendix is specific to each individual, but we drop the index i for expositional simplicity. For an unmarried individual, the social security benefit at age j is

$$\text{BEN}_j = \text{PIA}_j \times \text{CPI}_{62,j} ,$$

where PIA is the primary insurance amount, and $\text{CPI}_{62,j}$ is the cumulative inflation index from age 62 to the age at which the benefit is computed. Then the mortality-adjusted benefit is

$$E_{22}(\text{BEN}_j) = \text{BEN}_j \times P_j ,$$

where $E_{22}(\text{BEN}_j)$ is the expected value at age 22 of the benefit to be received at age j , and P_j is the conditional probability of survival to age j , given survival to age 22. For married individuals, the basic benefit is computed in the same manner. We compute the spousal benefit for the wife (or analogously, the husband) as

$$\text{SpBEN}_j = 0.5 \times \text{SBEN}_{j_s} ,$$

where SpBEN_j is the spousal benefit at wife's age j , SBEN_{j_s} is the husband's PIA adjusted for inflation to age j_s , and j_s is the husband's age when the wife is age j . Similarly, we calculate the survivor benefit:

$$\text{SurvBEN}_j = \text{SBEN}_{j_s} ,$$

where SurvBEN_j is the wife's survivor benefit after the death of the husband. If the other spouse is alive, we assume that a married individual receives the greater of his or her own benefit (BEN) and the spousal benefit (SpBEN). If the other spouse is deceased, the individual receives the greater of his or her own benefit (BEN) and the survivor benefit

(SurvBEN). Using PH_j and PW_j for the husband's and wife's survival probabilities, the husband's mortality-adjusted benefit is

$$E_{22}(\text{HBEN}_j) = PH_j[PW_j \max(\text{BEN}_j, \text{SpBEN}_j) + (1 - PW_j) \max(\text{BEN}_j, \text{SurvBEN}_j)],$$

where $E_{22}(\text{HBEN}_j)$ is the expected value at age 22 of the husband's benefit. This expected value includes only the dollars going directly to the husband. A symmetrical calculation is made to determine the wife's mortality-adjusted benefit:

$$E_{22}(\text{WBEN}_j) = PW_j [PH_j \max(\text{BEN}_j, \text{SpBEN}_j) + (1 - PH_j) \max(\text{BEN}_j, \text{SurvBEN}_j)].$$

We then compute the present value of expected taxes and benefits at age 22 for each individual, using alternative values for the constant real discount rate r :

$$\text{PVTAX} = \sum_j \frac{E_{22}(\text{SST}_j)}{(1+r)^{j-22}},$$

$$\text{PVBEN} = \sum_j \frac{E_{22}(\text{BEN}_j)}{(1+r)^{j-22}}.$$

Finally, the internal rate of return is computed by finding the discount rate that equates these present values.

REFERENCES

- Altig, David, Alan J. Auerbach, Laurence J. Kotlikoff, Kent A. Smetters, and Jan Walliser (1997). "Simulating U.S. Tax Reform." Washington, DC: Congressional Budget Office. Technical Paper Series.
- Auerbach, Alan J., and Laurence J. Kotlikoff (1987). *Dynamic Fiscal Policy*. New York: Cambridge University Press.
- 1994–1996 Advisory Council on Social Security (1997). *Findings and Recommendations*. Washington, DC: Social Security Administration.
- Boskin, Michael J., Laurence J. Kotlikoff, Douglas J. Puffert, and John B. Shoven (1987). "Social Security: A Financial Appraisal Across and Within Generations." *National Tax Journal* 40(March):19–34.
- Burkhauser, Richard V., and Jennifer Warlick. (1981). "Disentangling the Annuity from the Redistributive Aspects of Social Security." *The Review of Income and Wealth* 27(December):401–421.
- Caldwell, Steven, Melissa Favreault, Alla Gantman, Jagadeesh Gokhale, Thomas Johnson, and Laurence J. Kotlikoff (1999). "Social Security's Treatment of Post-war Americans." (This volume.)

- Coronado, Julia Lynn, Don Fullerton, and Thomas Glass (1998). "The Progressivity of Social Security." Department of Economics, University of Texas at Austin.
- Duggan, James E., Robert Gillingham, and John S. Greenlees (1993). "Returns Paid to Early Social Security Cohorts." *Contemporary Policy Issues* 11(October): 1-13.
- _____, _____, and _____ (1995). "Progressive Returns to Social Security? An Answer from Social Security Records." Washington, DC: U.S. Treasury Department, Office of Economic Policy. Research Paper no. 9501.
- Duleep, Harriett Orcutt (1986). "Measuring the Effect of Income on Adult Mortality Using the Longitudinal Administrative Record Data." *Journal of Human Resources* 21(Spring):238-251.
- Faber, Joseph F., and Alice H. Wade (1983). *Life Tables for the United States: 1900-2050*. Washington, DC: Social Security Administration, Office of the Actuary. Actuarial Study no. 89.
- Flowers, Marilyn, and John B. Horowitz (1993). "Distributional Implications of the Social Security Spouse Benefit." *Southern Economic Journal* 60(October):490-494.
- Fullerton, Don, and Diane Lim Rogers (1993). *Who Bears the Lifetime Tax Burden?* Washington, DC: The Brookings Institution.
- Garrett, Daniel M. (1995). "The Effects of Differential Mortality Rates on the Progressivity of Social Security." *Economic Inquiry* 7(July):457-475.
- Gramlich, Edward M. (1996). "Different Approaches for Dealing with Social Security." *Journal of Economic Perspectives* 10(Summer):55-66.
- Hamermesh, Daniel S., and Albert Rees (1993). *The Economics of Work and Pay*. New York: Harper Collins College Publishers.
- Heckman, James J., and Thomas E. MaCurdy (1980). "A Life Cycle Model of Female Labor Supply." *Review of Economic Studies* 47(January):47-74.
- Hurd, Michael D., and John B. Shoven (1985). "The Distributional Impact of Social Security." In *Pensions, Labor, and Individual Choice*, David A. Wise (ed.). Chicago: University of Chicago Press for NBER.
- Ibbotson Associates (1998). *Stocks, Bonds, Bells, and Inflation 1998 Yearbook*. Chicago: Ibbotson Associates.
- Jorgenson, Dale W., and Daniel T. Slesnick (1987). "Aggregate Consumer Behavior and Household Equivalence Scales." *Journal of Business and Economic Statistics* 5(April):219-232.
- Kotlikoff, Laurence J., Kent A. Smetters, and Jan Walliser (1998). "Social Security: Privatization and Progressivity." *American Economic Review* 88(May):137-141.
- Lee, Julie, Mark McClellan, and Jonathan Skinner (1999). "The Distributional Effects of Medicare." (This volume.)
- Myers, Robert J., and Bruce D. Schobel (1983). "A Money's-Worth Analysis of Social Security Retirement Benefits." *Society of Actuaries Transactions* 35(October):533-561.
- Panis, Constantijn W.A., and Lee A. Lillard (1996). "Socioeconomic Differentials in the Returns to Social Security." Santa Monica, CA: RAND Distribution Services. RAND Labor and Population Program Working Paper Series 96-05.
- Pechman, Joseph A., and Benjamin A. Okner (1974). *Who Bears the Tax Burden?* Washington, DC: The Brookings Institution.
- Rofman, Rafel Paulino (1993). "Social Security and Income Distribution: Mortality and Equity in Pension Plans." University of California, Berkeley. Ph.D. Dissertation (Demography).

- Rogot, Eugene, Paul D. Sorlie, Norman J. Johnson, and Catherine Schmitt (1992). *A Mortality Study of 1.3 Million Persons by Demographic, Social, and Economic Factors: 1979–1985 Follow-up*. National Institutes of Health. NIH Publication 92-3297.
- Steuerle, C. Eugene, and Jon M. Bakija (1994). *Retooling Social Security for the 21st Century: Right and Wrong Approaches to Reform*. Washington, DC: Urban Institute Press.
- U.S. Department of Health and Human Services (1993). *Vital Statistics of the United States—1989*. Hyattsville, MD: Public Health Service, Centers for Disease Control and Prevention, National Center for Health Statistics.
- U.S. Social Security Administration (1998). *Annual Statistical Supplement, 1997*. Social Security Bulletin. Washington DC: U.S. Government Printing Office.
- U.S. Social Security Administration (1998). *The 1998 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds*. Washington, DC: U.S. Government Printing Office.
- Williams, Robertson C. III (1998). "Equity in Social Security Benefit Determination: The Effect of Ignoring Some Working Years." Department of Economics, Stanford University.