Long Run Effects of Social Security Reform Proposals on Lifetime Progressivity

September, 1999

Julia Lynn Coronado
Federal Reserve Board of Governors
Washington, DC 20551
jcoronado@frb.gov

Don Fullerton
Department of Economics
University of Texas at Austin
Austin, TX 78712
and NBER
dfullert@eco.utexas.edu

and

Thomas Glass
Glass & Company, CPAs
515 Congress Ave., Suite 1900
Austin, TX 78701
tglass@glasscpa.com

This paper will be presented at the National Bureau of Economic Research (NBER) conference on "The Distributional Effects of Social Security Reform," Woodstock, VT, October 22, 1999. The model used in this paper reflects helpful comments from Angus Deaton, Martin Feldstein, Jeff Liebman, Gib Metcalf, Jim Poterba, and Jon Skinner. This paper is part of the NBER 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.

1. Introduction

Many observers believe that something must be done about Social Security. Although the original "pay as you go" (PAYGO) system was converted to a partially-funded system in 1983, promised future benefits still exceed expected future taxes -- especially by the time the "baby boom" population bulge is done with their retirement. When converted into 1995 dollars, the "intermediate" projected deficit for the year 2075 is $480 billion (Board of Trustees of the Social Security System, 1998). In addition, Social Security has been labelled "unfair." It redistributes not only from current working generations to the retired, but it also redistributes between families in different circumstances. Social Security provides benefits in excess of taxes to a spouse who stays at home and to women generally (because they tend to live longer than men). It penalizes work after the normal retirement age, and yet it continues to pay benefits to nonworking retirees who may be wealthy.

Because of these perceived problems, several fundamental reforms have been suggested. Some would eliminate the projected shortfall, switching to a fully-funded system, while others would eliminate the current trust fund and return to the original PAYGO idea. Some would eliminate the transfer elements of Social Security altogether, switching to mandatory individual defined-contribution accounts, while others would embrace the need for transfers and enhance them. After all, Social Security was meant to be part of a social "safety net," providing a larger replacement rate for low-income workers than for high-income workers.

In this paper, we estimate changes in progressivity from several particular reforms that have been suggested. We are not concerned with redistributions between family types, or between men and women, only redistributions between rich and poor. To define who is rich or poor, however, we use an estimate of *lifetime potential* income -- the present value of the total value of one's time. Thus a spouse who chooses to stay at home is not misclassified as "poor".

We use a large data set of almost two thousand individuals, and classify them into five lifetime income groups. For each person in each quintile, we calculate the present value of the Social Security tax they would pay and benefits they would receive. The difference is divided by lifetime income to provide a lifetime "net tax rate" for each income group. If this net tax rate rises across the five income groups, the system (or reform) is deemed progressive.

We cannot possibly evaluate all reforms that have been suggested, so we focus on four that are meant to represent the broad spectrum of possibilities. First, one set of reforms would privatize social security or switch to mandatory individual accounts with benefits that depend on contributions (e.g. Feldstein and Samwick, 1998). Since such a plan does not redistribute, but provides benefits equal to the present value of one's own contributions, the net tax rate is zero in our model. The effects in our model are the same as the repeal of Social Security, and we show those effects below. Second, we evaluate the proposal of the National Commission on Retirement Policy (NCRP, 1999). This plan redirects two percentage points of tax into defined-contribution individual accounts, and it drammatically cuts other benefits to balance the Social Security budget at that reduced tax rate. Third, we look at the plan of Aaron and Reischauer (1998), which suggests smaller specific changes to balance the long-run budget without fundamentally altering the nature of Social Security. Fourth, we calculate effects of the Moynihan (1999) plan that depletes the current social security trust fund through lower tax rates and then switches to "pay as you go."

The last three plans would speed up the currently-scheduled increase in the normal retirement age to 67, a change that affects transition generations. Furthermore, the NCRP and Moynihan plans would continue to phase-in retirement age increases to 70. Because we focus on redistributions between rich and poor, we do not calculate redistributions between young and old. Our model is designed only to look at the long-run provisions of social security and reforms.

Thus we capture the effects of the eventual change in the retirement age from 67 to 70, but we do not estimate effects on transition generations.

The next subsection overviews the model, and the following subsection overviews the existing literature on redistributive effects of Social Security. After this introduction, section 2 provides more detail on the model, and section 3 provides more detail on the four suggested reforms. Section 4 discusses our basic results, and section 5 discusses the sensitivity of those results to alternative assumptions. Section 6 concludes.

1.1 Overview of the Model

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.[1] Within this steady state context, we take account of the ways in which social security redistributes across groups defined by income, gender, and marital status. That is, while we report only the redistributions between lifetime income quintiles, we account for heterogeneity within each such quintile. Thus we capture the fact that different income groups have different proportions of individuals who are single or married, male or female, work continuously or sporadically, and who have different mortality rates.

We use twenty-two 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 secondary workers, and part-time secondary workers). The estimated coefficients are used to project each individual's wage rates before and after the sample period, so that each individual has

---

[1] The model used in this paper was introduced in Coronado et al (1998) to study the existing social security system, and it was used in Coronado et al (1999) to study generic reform components. Here, we evaluate specific reform packages.

a complete wage profile from age 22 to 66 (extended through 69 for plans with retirement at 70). 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.

Next, 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. We then derive income-differentiated mortality rates, and we use those mortality probabilities with constructed earnings histories to calculate each individual's expected lifetime social security taxes and benefits. Finally, we add over the individuals in each quintile to get the net impact of social security on each group.

Using actual earnings data is one of the important innovations of our model. 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 or complete privatization. In addition, we calculate effects of other specific reforms, and we compare the progressivity of those reforms to a proportional cut in all benefits (with a comparable overall net tax rate).

1.2 Overview of Existing Literature

The social security system takes taxes from both a high-wage person and a low-wage person during working years, and it provides benefits to both individuals when retired. We wish to measure how much of this money is transferred between individuals, rather than just transferred from the working years to the retirement years of the same person.

Initial tax incidence studies like Pechman and Okner (1974) used groupings based on annual income. This type of study would find that the social security system is progressive, but it aggregates unlike individuals. The low-annual-income group may include both the working poor and those who have retired from a high-earning career. Some later studies like Auerbach and Kotlikoff (1987) include lifetime profiles and lifetime decision making, in order to find how social security redistributes between young and old. However, this study does not distinguish between different lifetime income groups of the *same* age.[2]

Although much work has focused on intergenerational effects of the social security system, considerable work has also looked at intragenerational redistribution - using arbitrary levels of income for different groups. For example, Hurd and Shoven (1985) and Boskin et al (1987) each use three groups (e.g. median income, half the median, and five-times the median).[3] The approach of using arbitrarily-set income levels has tremendous computational appeal. However, the calculation of social security benefits depends not 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

[2] Nelissen (1998) finds substantial differences between annual incidence and lifetime incidence for social security in the Netherlands.

[3] Panis and Lillard (1996) set the low group at full-time minimum wage earnings, the middle group at the Social Security's 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).

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), Hurd and Shoven (1985), and Liebman (1998)]. Duggan, et al (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.[4] 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 incidence 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 et al (1998) 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 contain homogeneous individuals. Since everyone in a group must work the average amount for that group, these general equilibrium models cannot incorporate heterogeneity such as the fraction in each group who have zero earnings.

For these reasons, we do not attempt to build a general equilibrium model. The point of

---

[4] The true earnings can be estimated, however. For example, Fox (1982) uses information on the time of year that an individual reaches the wage cap to infer the full annual earnings.

this paper is to make use of much actual data on diverse individuals within each lifetime income group. In this way, we can look at distributional impacts of specific elements of the social security system.[5]

The literature on distributional impacts of specific elements of the social security system is small. Flowers and Horwitz (1993) look at the spousal benefit, whereby low-earner spouses can draw the greater of their own computed benefit or one-half of the higher-earning spouse's benefit. They demonstrate that the spousal benefit calculation is progressive when compared to an own-benefit calculation. This result is driven by their finding that higher-income families consist of spouses with more-equal earnings and lower-income couples have more disparate earnings. Our data imply the opposite: more-equal earnings among couples with equally-low wages. 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 under-reported in the social security records, especially for working-age individuals and minorities. Garrett (1995) uses mortality estimates from a literature search, while Panis and Lillard (1996) extract mortality information from the PSID. Since high-income people live longer, several studies show that accounting for income-differentiated mortality seriously dampens the progressivity of social security (e.g. Steuerle and Bakija, 1994, Duggan et al, 1995, and Panis and Lillard, 1996).

---

[5] 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 et al (1999) calculate such effects for Medicare.

Finally, Caldwell et al (1999) use a large micro-simulation model to construct lifetime earnings for many heterogeneous individuals. This model starts with the 1960 Census Public-Use Microdata Sample and uses estimated transition probabilities to "grow" the sample in one 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, post-war generations receive smaller and even negative rates of return.

2. Lifetime Earnings Profiles and Net Taxes from Social Security

In this section we describe the data and methodology used to obtain lifetime earnings profiles, to estimate mortality probabilities that differ by lifetime income, and to calculate net taxes from social security. A more detailed description is provided in the Appendix. We use the PSID for the years 1968 to 1989, which gives us twenty-two years of actual earnings data for a sample of the population. We select a sample consisting of 1086 heads and 700 wives that is 66 percent of the representative cross-section. This use of a reduced sample suggests the possibility of bias in our econometric estimates and our conclusions about the progressivity of social security. We do not believe our results are biased, however, for reasons discussed in Appendix A1.

The PSID provides only twenty-two years of actual data. In order to obtain complete profiles of earnings from age 22 through age 66 for each of our sample members, we want to be able to generate out-of-sample earning observations.[6] We do this by estimating earnings

[6] We assume that people work until the future normal retirement age of 67, claim social security benefits at that point, and do not work after retirement. While the majority of people retiring in the past decade have claimed early retirement, they get a reduction in benefits that is (continued...)

regressions and using the estimated coefficients to generate the needed observations. However, as Fullerton and Rogers (1993) demonstrated using data from the PSID, earnings profiles can have significantly different shapes for different lifetime income groups. We therefore estimate separate earnings regressions for different lifetime income classes.

Our model is somewhat stylized in that we ignore inheritances and transfers. Our measure of annual income is based on wages, which are zero for a retired person. Lifetime income is the present value of that annual income. Note that capital income from lifecycle savings is *not* part of lifetime income. If the present value of consumption must equal the present value of labor income, then capital income just reflects rearrangements in the timing of consumption.

Then, to each lifetime earnings profile, we apply mortality probabilities that vary by our measure of potential lifetime income, gender, and race. We calculate expected social security taxes paid for each person for each year of their working life and expected benefits received during retirement, following current provisions of the Social Security Administration.

2.1 Lifetime Income

We want to estimate a separate earnings regression for each lifetime income class, and we want a measure of lifetime income that accurately reflects economic well-being. To begin, we calculate an annual wage rate for each member of our sample by dividing annual earnings by hours worked. To construct a wage rate for every year of each sample member's working life, we first use all positive wage observations to estimate log wage profiles. We estimate separate log wage

---

(...continued)
supposed to be actuarially fair. Early retirees have less education and are more likely to be retiring from blue-collar jobs, so they may have lower lifetime incomes (U.S. Congressional Budget Office, 1999). As discussed below, people with fewer lifetime resources have higher mortality rates, so the reduction in benefits with early retirement is likely to be too great.

regressions for heads, full-time working wives, and part-time working wives. The results of these regressions can be found in Appendix A2. We regress the log of the wage rate on an individual fixed effect and other variables like age, age-squared, and age-cubed. 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. Using the resulting fixed effects and coefficients, we then fill in missing observations during the sample period and observations outside the sample period. The appendix details how we assign a wage rate to women who have no earnings histories. Non-working wives do engage in household production, and assigning them a zero wage may incorrectly place them in a low lifetime income group for the earnings regressions and for the distributional analysis. Thus, for each individual, we have a wage rate for every year of their entire economic life from age 22 to 66.

We then use this wage rate and multiply it in each year by 4000 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.[7] Using this endowment allows us to abstract from the actual labor/leisure choice, since someone who chooses to work less and 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.[8]

Once we have a complete wage profile for each of our heads and wives for ages 22-66, we

---

[7] For sensitivity analysis, we show net tax rates with two other measures of income: the present value of actual earnings, and the present value of potential earnings where leisure is valued at the average wage rate for the sample instead of the individual's wage rate.

[8] On the other hand, our model may overstate the value of time at home to the extent that it represents sick days or involuntary unemployment.

calculate individual gross lifetime income as:

$$LI = \sum_{t=1}^{45} [(w_t \times 4000)/(1+r)^{t-1}] \tag{1}$$

where t indexes the forty-five years in the individual's economic lifetime relevant for social security, ages 22 to 66, and where the individual could work a maximum of 80 hours per week for 50 weeks per year. Through most of our analysis, we use a value of 2% for r, the real discount rate. Later, we see the effect of changing the discount rate.

As couples generally pool their resources, it would be inappropriate to place husbands and wives individually into separate lifetime income groups. The low-wage wife of a high-wage husband is not "poor". We therefore combine the lifetime income of the husband and wife, and divide by two, to obtain individual lifetime income for each of them. We can now deal with each member of our sample as an individual and categorize them into five lifetime income groups. The first quintile has the lowest income, and quintile five has the highest income.

2.2 Earnings Profiles

Once we have people classified into lifetime income groups based on what we feel to be an appropriate measure of economic well-being, we estimate regressions for actual earnings. For each quintile, using our data from the PSID, Appendix A3 describes how we estimate separate earnings regressions for heads, working wives, and occasional working wives, for a total of fifteen regressions. We use both positive and zero earnings observations in a Tobit framework.

Because the Tobit framework is nonlinear, we do not include fixed effects as their inclusion would imply inconsistent parameter estimates. The exclusion of fixed effects also means we can use variables in these earnings regressions that do not vary over time, such as education, race, and gender. For each regression for the heads, we begin with independent variables for age, age-squared, age-cubed, education, education-squared, the product of age and education, a dummy for whether the head is female, age interacted with the female dummy, and a dummy for whether the head is white. We then eliminate the variables that were insignificant. We follow a similar procedure for habitually working wives and part-time working wives.

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. We do not use these coefficients to fill in missing earnings observations during the sample period, as 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. In fact, we also simulate a representative number of zero earnings years for the out-of-sample portions of each earnings profile.

2.3 Income-Differentiated Mortality

It is a stylized fact that people with higher lifetime incomes tend to live longer, a fact that can dampen the progressivity of the benefit structure of the social security system. We derive a set of mortality probabilities that vary by race, gender, and our measure of potential lifetime income, so that we can examine the impact of differential mortality on redistribution. Standard mortality tables extend only to age 85 and are differentiated only by sex and race. As described in Appendix A4, we extend these data in three ways. First, we 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 present 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). For 100,000 individuals alive at age 0, these tables show the number surviving at each age 1 through 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. 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%, survival to age 65 of 75.82%, and survival to age 85 of 22.34%. 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 present value of 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 other 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, nonwhite males, and nonwhite females). Since 31% of the population is still alive at age 85, Appendix A4 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 being alive in each year from age 22 to 99.

Many studies have noted that mortality rates for the poor are larger than average. *A Mortality Study of 1.3 Million Person* (Rogot, et al, 1992) provides a rich source of data on this effect. They show the observed number of deaths for each annual income class of each race, gender, and ten-year age group. For each such cell, we divide observed deaths, O, by the expected deaths, E, that would occur if all income classes of that group had the same mortality rate. We then apply that O/E ratio to each cell in the extended mortality tables.[9] Among white males aged 25-34, for example, those in the poorest annual income group die at a rate that is 168% of the average, while those in the richest annual income group die at rate that is only 61% of the average. For nonwhite females of the same age, the poor die at a rate that is 186% of the average, while the rich die at a rate equal to 44% of the average.

Although we have the annual household income of each individual in our sample for each year, we do not just use the corresponding annual income group's O/E ratio for that person in that year to weight their mortality probability. Using annual income would imply that an individual with a steeply hump-shaped earnings profile would have a probability of dying that falls dramatically during high-annual-income years and then rises 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. To address this issue, our procedure described in Appendix A4 is based on the relative ranking of each individual's lifetime income. Basically, a person in a particular percentile of the lifetime income distribution is assigned the O/E ratio of a person in the same percentile of the annual income distribution.[10]

2.4 Social Security Taxes Paid

We next compute the value of social security taxes for each person in each year, following

[9] Income-differentiated mortality rates are also employed by Poterba (1997), Caldwell et al (1999) and Lee et al (1999), and Sabelhaus and Pence (1998).

[10] 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.

the provisions of the Social Security Administration. This tax is commonly called 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 separates trust funds, and benefits are paid from the appropriate fund. The program has become almost universal -- 95% of all employment in the U.S. is covered.[11]

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

Since an objective of our research is to measure each worker's net social security tax burden, the question arises: how much of the total FICA tax does the worker bear? Using only the statutory incidence (the worker's half) would yield much lower burdens 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

[11] 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.

employee tax.[12]

Our focus is the retirement portion of the social security system, not the disability insurance or hospital insurance. Of the total 15.3% tax rate, 2.9% is for Medicare (HI), leaving 12.4% for OASDI. This is the rate cited and modified by certain reform proposals, even though 1.8% goes to Disability Insurance (DI). The remaining 10.6% is for Old Age and Survivors Insurance (OASI), and this is the tax in our model.[13] The OASI portion of the tax is paid directly to the OASI Trust Fund, which is used to pay all retirement benefits. We ignore the DI and HI portions of the tax, as well as benefits paid from the DI and HI Trust Funds.

Our sample from the PSID includes observed and constructed earnings for each individual from age 22 through 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 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}):

$$N_{fij}=N_{oij}(AE_{fj}/AE_{oj}) \quad . \tag{2}$$

[12] Panis 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 Panis and Lillard, however, and for comparability with other studies, we treat the entire 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.

[13] These allocation percentages are for the year 2000 and beyond. Congress "temporarily" increased the portion going to DI for the years 1994 to 1996, followed by a reduction for 1997-1999. The 1997 allocation is: OASI = 10.7%, DI% = 1.7%, and HI = 2.9%.

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.[14]

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 individual i 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) as the number in cell i alive at age j divided by the number in cell i alive at age 22. We then calculate $E(SST_{ij})$, the expected social security tax of person i in year j, as:

$$E(SST_{ij}) = [T \times \text{Min}(N_{ij}, CAP_j)] \times P_{ij} \qquad (3)$$

where T is the combined OASI tax rate (which is constant with unchanged law), CAP_j is the maximum nominal earnings subject to the OASI tax (which increases with inflation), and P_{ij} is the probability that person i is alive at age j. These amounts are used to compute the present value of social security taxes paid.

2.5 Social Security Benefits

Under provisions of the Social Security Act, benefits are calculated from a progressive formula based on 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. Only earnings at or below the taxable cap in each year are considered.

---

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

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 works from age 22 through age 66 (retiring on his or her 67th birthday) would have a total of 45 years of earnings. Under the Act, only the highest 35 years are considered, so the ten lowest years will be dropped. AIME is the simple average of the indexed earnings in those 35 highest-earnings years.[15]

Next, the Primary Insurance Amount (PIA) is calculated as 90% of AIME up to the first bend point, plus 32% of AIME in excess of the first bend point but less than the second bend point, plus 15% of AIME in excess of that second bend point. The fact that only capped earnings are used to calculate AIME provides a *de facto* maximum benefit. In 1995, the bend points were \$426 and \$2,567. If AIME were \$3,200, for example, the PIA would be:

$$PIA = .90\times(426) + .32\times(2{,}567 - 426) + .15\times(3{,}200 - 2{,}567) = \$1{,}163.47 \quad (4)$$

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

A retiree is entitled to a benefit equal to the Primary Insurance Amount upon normal retirement at age 67. A worker may still choose to retire as early as age 62, with reduced benefits.[16] In contrast, if a worker elects to *delay* receipt of benefits to an age as late as 70, the

---

[15] The language of the Act specifies dropping the *five* lowest years of earnings through age 61. Then, if the worker has years of earnings after age 61 that are higher than some earlier years' earnings, the higher post-61 earnings will replace those lower earnings. The net effect for a worker retiring at age 67 is to drop the ten lowest years.

[16] This early retirement penalty is a permanent reduction in the PIA of 5/9% for each early month (6.67% for each early year). For example, a worker retiring at age 64 when the normal (continued...)

eventual benefits are permanently increased by 5% per year of delay. 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] which for our hypothetical cohort is 67.

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. The spouse of a retired worker can receive the greater of the benefit based on his or her own earnings, or one-half of the PIA of the retired worker (designated as the "spousal benefit"). 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 higher of the benefit based on his or her own earnings, or 100% of the benefit to which that worker was entitled. The benefit based on the deceased worker's benefit is called the "survivor benefit". We ignore non-spousal survivor benefits, which in aggregate are relatively minor.[18]

Our calculations of these amounts are detailed in Appendix A5. We use each individual's observed and constructed earnings to compute Average Indexed Monthly Earnings (AIME), the Primary Insurance Amount (PIA), the Spousal Benefit (SpBen), and the Survivor Benefit for the surviving spouse (SurvBen) in exact accordance with provisions of the Act.

2.6 Present-Value Net Tax Rates

(...continued)
retirement age is 67 would receive a benefit for the rest of his or her life that is reduced by 20%.

[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 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 to retired workers or their spouses, and only $14.8 billion (4.9%) was paid for the other survivor and miscellaneous benefits (*Annual Statistical Supplement, 1997,* Table 4A.5).

After we calculate the mortality-adjusted tax and benefit in each year for each individual in each of our lifetime income quintiles, we compute the present value, at age 22, of the benefits to be received minus the taxes paid. We then add over the individuals in each lifetime income quintile. This difference indicates the absolute size of the social security transfers between income groups. We divide by the present value at age 22 of the lifetime endowment (discounted at the same rate) to get an effective net tax rate for each group. A system that takes the exact same fraction of income for all groups is "proportional", whereas a system that takes a higher fraction of the income of the rich (poor) is deemed progressive (regressive).

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% for their long-term estimate of real returns.[19] 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%, and the nominal return on intermediate-term U.S. Treasury obligations was 5.4%, so the real rate of return was 1.4%.

For one choice of discount rate we use 2%, which lies between the forecast rate earned by the OASI trust fund on its investments (2.8%) and the historical average of real returns on government bonds reported by Ibbotson (1.4%).[20] To test the sensitivity of results, we also use

---

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

[20] Other studies of social security redistribution have used rates on either side of 2%. Myers and Schobel (1983) use 2%, Hurd and Shoven (1985) use 3%, Boskin, et al (1987) use 3%, Duggan, et al (1993) use 1.2%, Steuerle and Bakija (1994) use 2%, and Gramlich (1996) uses 2.3%. In contrast, Caldwell et al (1999) use 3%, 5%, or 7%.

a discount rate of 4%. As shown below, the choice of rate impacts not only the absolute size of the present value gains or loss for each group but also the pattern of progressivity.

3. Proposed Reforms and their Treatment in our Model

Our evaluation of social security reform is limited in many respects. First, since we focus on distributional effects, we ignore behavioral effects such as changes in labor supply or savings. Second, since we cannot evaluate all of the many suggested reform proposals, we focus on only four of the major ones. Third, since most of these proposals are still evolving, we evaluate only versions that were available in written form in early 1999. Fourth, since each such proposal is too complicated to capture fully in our model, we really just evaluate a "stylized" version of each reform. In particular, since we consider only long run provisions, we ignore any phase-in of a change in the normal retirement age.[21] Since we assume everybody retires at that normal retirement age, we also ignore the effects of proposed changes in the early retirement age. And since we have only one "discount rate" in our model, with no consideration of risk, we cannot capture the welfare effects of any plan to switch some of the social security trust fund from government bonds to investments in corporate stocks and bonds.

Because we miss some of these ways in which each reform might raise net revenue, especially during the transition, we cannot comment on the extent to which each reform might close the existing social security deficit. Perhaps all the plans balance that budget in present value, but the long-run features that we consider raise different amounts of net revenue for each

[21] 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 2017, resulting in a normal retirement age of 67 after the year 2021. All of the reform proposals we consider would eliminate the pause (from 2005 to 2017) and reach the new higher normal retirement age sooner.

plan. As a consequence, some of the plans appear in our model to have higher overall net tax rates than others. We emphasize, however, that our goal is to compare the progressivity of these plans and not their overall net tax rates.

As described above, we do capture the major long-run provisions of social security that determine taxes and benefits for individuals in different circumstances. We now describe proposed changes to these provisions, as summarized in **Table 1**. The first column of this table represents the current system. It does not list all features of the current system, only the main ones that would be reformed by one of the plans.

3.1 The Feldstein and Samwick (F&S) Plan

A number of proposals would completely privatize social security. The proposal outlined by Feldstein and Samwick (1998) is typical of these plans. It specifies a transition from the current system to one in which the benefits are equivalent to those guaranteed under the current system, but where these benefits in the long run are funded entirely by mandatory contributions to individual accounts made over a lifetime. The balances in the individual accounts would be invested in private debt and equity markets.

We do not explicitly model the Feldstein-Samwick plan, or any other plan based on individual accounts, as such plans involve little or no redistribution (except to the extent that some privatization plans include minimum benefits and survivor benefits). In our modeling framework, such plans are equivalent to repeal of the system. Our model is better suited to capturing the effects of reforms that alter the basic tax and benefit nature of the current system. Thus, in our model, the effects of the Feldstein-Samwick plan are the opposite of the effects of having the current social security system.

3.2 The National Commission on Retirement Policy (NCRP) Plan

The NCRP (1999) proposal is also associated with the names of Senators Breaux and Gregg and Representatives Kolbe and Stenholm. One version is a "defined benefit" (DB) plan based on the current OASDI tax rate, but we evaluate only the other version that sets aside 2 percentage points of each person's tax into a mandatory "Individual Savings Account" (ISA). Since each retired individual receives back his or her own ISA, plus earnings, such a mandatory savings plan does not redistribute. It has a net present value tax of zero for each individual and therefore does not enter our calculations. The remaining "tax and benefit" portion of this plan is scaled back from current law. As shown in the first row of Table 1, the OASDI tax rate is cut from 12.4% to 10.4%.[22] The next row shows that this plan does not change the wage cap.

With taxes cut substantially, closing the social security deficit means that benefits must be cut dramatically. The NCRP plan cuts benefits in several ways. The largest cut is in the calculation of benefits called the "primary insurance amount" (PIA) in equation (4). In that calculation, "PIA factors" are applied to average indexed monthly earnings (AIME) between the bend points. As shown in the third row of Table 1, the long-run NCRP plan would still provide 90% of AIME up to the first bend point, but the 32% rate between the two bend points is cut to 21.36%, and the 15% rate above the second bend point is cut to 10.01%. Thus benefits are cut disproportionately for high-income individuals. This change is progressive, even in a lifetime context, as we confirm below.

In fact, this plan adds benefits to low-income individuals, another progressive change. Current law has a small "minimum benefit" that depends on the number of quarters of earnings but

---

[22] These reforms state changes in terms of the current 12.4% OASDI rate, as shown in the Table. However, 1.8 percentage points of that tax go to disability insurance (DI), and we model only OASI. With the 2 points diverted into ISA accounts, the 10.6% OASI rate becomes 8.6%.

that can reach as high as $6,235 per year (at 1995 levels, but indexed). As indicated in the next row of Table 1, the NCRP plan would raise this minimum benefit to the indexed "aged individual poverty level" (AIPL), which was $7,761 in 1995 (a 24% increase).

The next biggest cut in benefits occurs through the NCRP's increase in the ultimate normal retirement age (NRA) from 67 to seventy.[23] This change is regressive for three reasons. First, it means that individuals will work and pay taxes for more years, and those taxes are generally regressive because they apply only to earnings below the wage cap. Second, it means that individuals will retire later and receive benefits for fewer years. Since the benefit schedule is progressive, that cut in benefits is regressive. Third, because of income-differentiated mortality, the higher retirement age means that low-income individuals have a disproportionate increase in the chance of dying before they receive retirement benefits.

The NCRP plan also changes the number of years' earnings used in the AIME calculation. Long-run current rules use 35 years, which means that ten years of low earnings can be dropped from the calculation. This includes the five years that can be dropped before the AIME calculation at age 62, and the additional five years of earnings from age 62 until retirement at age 67 that can be used to replace lower earnings from before 62. The NCRP plan says it would "include earnings for all years", and we interpret this to mean all years up to the AIME calculation. Since the NCRP plan raises the early retirement age (and timing of the AIME calculation) from 62 to 65, the individual can still use five subsequent years of earnings (until retirement at age 70) to replace lower earnings from before 65. In other words, effectively, 5

[23] Like other reforms considered here, the NCRP plan would also later increase the NRA *above* 70 to account for subsequent increases in longevity (to maintain a constant number of expected years of life after retirement. We cannot model this provision as an increase in the retirement age, unless we were also to raise survival probabilities (which would roughly *maintain* the expected number of years of benefits).

years still can be dropped.[24]

Under current law, any married retiree can receive the higher of one's own benefit or half of what the spouse receives. This latter option is called the "spousal benefit." The next row of Table 1 shows that the NCRP plan would allow only one-third of the spouse's benefits. This cut would most affect any person whose income is low relative to his or her spouse, but remember that we do not count that person as "poor" (because we assume each person gets half of the couple's total income). Perhaps surprisingly, this change is slightly progressive. As it turns out, middle-income couples have more disparate incomes and make greater use of the spousal benefit.

While the current system is fully indexed for inflation, it uses the consumer price index (CPI). This index has been criticized for overstating inflation and therefore scaling up benefits by more than necessary to maintain living standards for retired beneficiaries. The NCRP plan, like other reform plans, would require a downward revision in the CPI, which would raise some net revenue. If the issue was just described in terms of accurate indexation for inflation, then we would not be able to capture this provision. If the Bureau of Labor Statistics does not change the CPI, however, the reform says that benefits will be indexed explicitly to the CPI minus 0.5 percentage points. We model that change as a real cut in benefits. Specifically, real benefits fall at 0.5% per year, starting at the age of retirement. Since the benefit schedule is progressive, any cut in benefits would normally be regressive. However, benefits are cut more for those who live longer and continue to experience real benefit cuts each year. Since high-income individuals live longer, this particular form of benefit cut has uncertain effects. As it turns out, the net effect of

[24] This reform component was analyzed in Coronado, Fullerton, and Glass (1999), using the same model as in this paper, and it was found to be slightly regressive. Including the low earning years reduces AIME in a similar fashion for all income groups, but that decline in AIME reduces benefits at the 90% PIA-factor for individuals below the first bend point, and it reduces benefits at a low PIA-factor for those with income above the bend points.

this provision is somewhat regressive in our model.

Thus, some aspects of the NCRP plan are progressive, and some are regressive. Our calculations below will show the net effects of all these changes together.

Table 1 also lists a few provisions that are not captured by our model. The NCRP plan would also raise the age for early retirement from 62 to 65 (to match the three-year increase in the normal retirement age from 67 to 70). It would extend OASDI coverage to all state and local government employees hired after 1999.

Under current law, if a social security beneficiary works after normal retirement age, retirement benefits are reduced by $1 for every $3 that are earned above some threshold. This feature is not captured in our model, because we assume no earnings after retirement in any case. The NCRP plan would eliminate this retirement earnings test for individuals after normal retirement age (70).

3.3 The Aaron and Reischauer (A&R) Plan

Any reform plan must face fundamental choices about the very nature of social security. The current system is "partially funded," so a reform could raise revenue and create a fully funded program, or it could return to the original "pay as you go" (PAYGO) idea. The current system is explicitly a transfer program that redistributes from workers to retirees, to those with low income, to non-earning spouses, and to women (because they live longer). Any reform could choose to remove those transfer elements, or it could choose to enhance them.

Rather than make wholesale changes to social security, the plan of Aaron and Reischauer (1998) would "fix" the current system. It would "close the projected long-term deficit and make Social Security better reflect current social and economic conditions, while preserving Social Security's fundamental character" (p. 96).

As a consequence, this plan tinkers with a number of provisions in ways that each raise a bit of revenue. This plan is summarized in the third column of Table 1. As it turns out, many of these changes appear at the bottom of the column, under features "not captured by our model." The A&R plan would raise the earliest eligibility age (EEA) from 62 to 64 (to match the currently provided two-year increase in the normal retirement age from 65 to 67). Like other plans, A&R would cover all new state and local employees.

Whereas current law collects income tax on social security benefits only above some threshold, the A&R plan would tax social security benefits just like a private pension. That means it would exempt the amount that was already subjected to income tax (such as the employee's payroll tax share, which comes out of taxable income), but it would tax the rest of social security benefits -- since those dollars have not yet been subject to income tax. We model effects of social security only, ignoring the income tax, so we do not capture these interactions (see footnote 12).

The A&R plan would also raise some money by switching part of the Social Security trust fund from government bonds to higher-yielding corporate stocks and bonds. We use only one discount rate, ignoring different risk premia, so we do not capture this provision either. We might note, however, that many of these ignored provisions are unlikely to have major direct effects on progressivity.

The top of the third column shows the provisions of the A&R plan that *are* captured in our model. First, this plan would change the number of years of earnings used in the AIME calculation from 35 to 38. Including more low-earning years means that AIME is reduced, and thus benefits are lower. The calculation still drops four years before the AIME calculation, and it still uses three more years (from 64 to 67) to substitute for earlier low-earning years. Thus it drops the 7 lowest earning years before age 67.

Like the NCRP plan, the A&R plan would raise money by cutting the spousal benefit from one-half to one-third of the benefits of the higher-earning spouse. As mentioned above, the reduction in the number of dropped years is slightly regressive, and the reduction in the spousal benefit is somewhat progressive.

Next, the A&R plan makes a change to the "survivor's benefit," which currently allows a widow or widower to receive his or her own benefit or the deceased spouse's benefit, whichever is larger. In the table, this rule is represented by "Max(Hus, Wife)". Instead, the A&R plan would provide three-quarters of the *combined* benefits of both spouses ["3/4(Hus+Wife)"]. The logic for this change is based on the cost-of-living for one person compared to the cost for two together. Compared to current law, however, it provides more benefits to some individuals and less to others. If two spouses had the same earnings, for example, then either person's survival benefit would become three-quarters of the total, which is 50% *more* than under current law (where either person would get half of the total). If a lower-earning spouse had own benefits of less than 1/3 of the higher-earning spouse, then either person's new survivor's benefit would be *less* than under current law.[25] In our calculations, this particular provision is found to be progressive. As mentioned above, low-income couples tend to have more similar incomes, since both must work at low-paying jobs. Equal incomes gain from this reform provision. Middle and high-income couples tend to have more disparate incomes, since they can afford for one person to stay at home, and thus gain less or actually lose from this proposal.

Finally, the A&R plan would undertake unspecified corrections in the consumer price index (CPI). The reasoning is the same as that described above, namely, that the current CPI has been criticized for growing too fast. This plan would leave those corrections to the economics experts, however, and not subtract any number of points from the CPI. With the system fully

---

[25] The breakeven point is where one's benefit is a third of the higher earning spouse's benefit, because 3/4 of (one plus one-third) equals max(one-third, one).

indexed to an accurate measure of inflation, we assume that real benefits are maintained.

Again, some of these provisions are progressive and some are regressive. Most are small, however, and so the overall progressivity of the A&R plan is not expected to differ much from that of current law. As shown below, the A&R plan is slightly more progressive than the current Social Security system.

3.4 The Moynihan Plan

In terms of fundamental choices about the nature of Social Security, Senator Moynihan's (1999) reform proposal would head in a different direction. Whereas the 1983 changes raised revenue to generate a partially-funded Social Security trust fund, this plan would return to "pay as you go" (PAYGO). The current trust fund would be drawn down by a temporary *reduction* in the current 12.4% OASDI tax rate to 11.4% (for 1999-2000) and to 10.4% (for 2001-2024). Then, when the trust fund is depleted, and that tax on a smaller number of workers is not enough to cover the benefits for a larger number of retirees, the rate would have to rise again to 11.4% (for 2025-29), 12.4% (for 2030-44), 12.7% (for 2045-54), 13.0% (2055-59), and 13.4% thereafter. These numbers are summarized in the top of the fourth column of Table 1.

Since our model considers only the long-run provisions of these reforms, the Moynihan Plan has to be represented by the 13.4% tax rate. We show results with the 13.4% rate in our tables below. In the long run, with this rate, overall net tax rates on all individuals are substantially higher than for the other reforms (and higher than for current law). The reason is that this plan disperses the trust fund to those of us in current generations -- by lowering *our* overall net tax rates. For this reason, results below *also* show the effects of the Moynihan Plan with the low 10.4% rate (which applies in 2001-2024).

The wage cap was $68,400 in 1998, and it is projected to reach $82,800 in 2003. As indicated in the second row of the table under the current system, this wage cap will cover about 85% of wages. This percentage has been falling, because high wage rates have been growing faster than average wages. The Moynihan Plan would raise the wage cap to $97,500 in 2003, which would cover about 87% of wages, and it would still be indexed thereafter. We calculate the real increase in the long-run wage cap for our model. This change is progressive, because it collects additional payroll tax from those above the current wage cap. On the other hand, the increase in the OASDI tax rate to 13.4% is regressive, given any wage cap, because it collects only from those below the cap. Again, our model calculates the net effect on progressivity.

The Moynihan plan also speeds up the currently-scheduled increase in the normal retirement age to 67, and it continues that increase to the age of 70 (for those retiring in 2073 and later). This change is regressive, for three reasons mentioned above: first, individuals pay the regressive payroll tax for more years, and second, they receive the progressive benefit schedule for fewer years. Third, low-income workers die sooner, so the fall in their survival probability from age 67 to 70 is greater than for high-income workers.

Since it increases the normal retirement age by three years (from 67 to 70), the Moynihan Plan also increases the number of years of earnings in the AIME calculation by three (from 35 to 38). The lowest ten years of earnings are still ignored. This change is slightly regressive, for reasons mentioned in footnote 24.

Like the NCRP, the Moynihan plan requires a reduction in the index used to maintain real benefits after retirement. If those corrections are not made within the CPI, then benefits will be indexed by the CPI minus one percentage point. We model this change as a one percent cut in real benefits each year after retirement.

The bottom of the last column of the table indicates the provisions of the Moynihan Plan that are not captured by our model. Like other reforms, it extends coverage to all new state and local workers. Like the NCRP plan, it eliminates the current earnings test for those beyond the retirement age who work. While the NCRP plan would eliminate this test at age 70, the Moynihan plan would eliminate it at age 62. Finally, like the A&R Plan, the Moynihan plan would change the income tax to cover all Social Security benefits as if they were private pensions.

4. Results

Our initial simulations use the enacted provisions of the Social Security Act, applied to a future cohort born in 1990. Results are presented in **Table 2**. The first row shows the overall average *un*discounted taxes paid minus benefits received, in thousands of dollars per person (for the current system and each proposed reform). The reason for showing undiscounted net taxes 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 number of entering 22-year-olds in each of the sex-race-income cells in our model, the undiscounted sum of taxes paid per individual ($103,200) equals the sum of taxes paid by all ages alive at one time. Similarly the undiscounted sum of benefits ($164,900) is the sum of benefits paid out to all ages alive at one time. On this basis, the current social security system loses the difference ($61,700 per 22-year-old) each year.[26] Each proposal closes this shortfall to a greater or lesser degree, which provides an imperfect measure of comparability in our stylized model.

The first column of Table 2 shows the net tax as a fraction of lifetime potential income for

---

[26] If we multiply this $61,700 figure by the number of 22-year-olds alive in 1994 (about 3.7 million), we get a total loss of about $228 billion per year. This figure lies between the "low" and the "high" deficit projected by the Board of Trustees of the Social Security System (1998). As mentioned above, their "intermediate" projected deficit for 2075 is $480 billion in 1995 dollars, but that includes DI and pertains to a larger population.

each quintile under the current system. This net tax rate rises from 0.62% for the lowest-income quintile to 1.01% for the highest-income quintile. Thus current law is progressive, but not uniformly progressive. The highest net tax rate applies to the middle-income quintile (1.07%). The benefit structure is progressive, even on a lifetime basis, but that progressivity is largely offset by the regressive tax system (which exempts earnings above the wage cap) and by various features of the system that tend to favor high-income groups (like the fact that the rich tend to live longer and therefore receive benefits longer).

A large number of recent articles on social security reform have dealt with privatization of the system or other large-scale overhauls of the program (e.g. Kotlikoff et al, 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. Thus, the results in this first column for the current system can be viewed as the distributional impact of an extreme reform -- repeal of social security. Because the current system is progressive, its repeal would be a regressive change.

Feldstein and Samwick (F&S, 1998) do not suggest repeal of the current system, but instead outline a plan to make it solvent and actuarially fair. Most importantly, their plan deals with the costs of a transition that honors the currently-promised benefits to existing generations. That transition does not show up in our long-run model. If the F&S plan is actuarially fair, in the long run, then it does not take net taxes from anybody. The second column of Table 2 shows these zero net tax rates in the top panel, and the change from current law in the bottom panel. To the extent that our calculations accurately reflect the long-run progressivity of the current system, the change to zero tax rates would be a regressive change.

The results in Table 2 are illustrated in **Figure 1**, where the net tax rate for the current system is the lowest of these six lines, rising from 0.62% for the first income group to 1.07% for the middle group and falling to 1.01% for the high income group. The Feldstein-Samwick plan would be represented by the horizontal axis, with zero tax rates for all groups.

The NCRP plan has both progressive and regressive elements. It would lower the regressive payroll tax by 2 percentage points and make the benefit schedule more progressive, but it would also cut benefits by raising the retirement age, by including more years of earnings in the benefit calculation, and by reducing the CPI by 0.5 percentage points. Since benefits generally are progressive, these benefit cuts are regressive. The net effects of all these changes are shown for the NCRP plan in the third column of Table 2. The net tax rate on the lowest income group would rise to 1.83%, but the rate on the highest income group would rise to only 1.77%. Again the middle group pays the highest tax rate (2.15%). The bottom panel of Table 2 shows that the increase in tax rate is highest for the group whose income is lowest. Thus the reform is a regressive change to the current system. On the other hand, Figure 1 shows that the NCRP system overall is pretty flat, with a net tax rate around 2% of lifetime income for all groups.

The top of the NCRP column in Table 2 shows the annual shortfall. By raising the net tax rate for everybody, the NCRP plan is able to reduce the annual shortfall as measured in our model from $61,700 per person to only $5,600 per person (and provisions that we do not capture may raise the rest of the needed revenue to balance the social security budget). Yet, one might ask, if "balance" means that all social security benefit payments are covered by payroll taxes, then why is the net tax rate still positive (at about 2% of lifetime income)? A zero balance in our model is represented by total taxes equal to benefits without discounting, to represent all cohorts alive at one time. In contrast, the net tax rate in our model is the discounted present value of one's taxes minus one's own benefits during life. Since taxes come before benefits, discounting means that