

NBER WORKING PAPER SERIES

IS SOCIAL SECURITY PART OF THE SOCIAL SAFETY NET?

Jeffrey R. Brown
Julia Lynn Coronado
Don Fullerton

Working Paper 15070
<http://www.nber.org/papers/w15070>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

This research was supported by the U.S. Social Security Administration through grant #10-P-98363-1-03 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Jeffrey R. Brown, Julia Lynn Coronado, and Don Fullerton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Is Social Security Part of the Social Safety Net?
Jeffrey R. Brown, Julia Lynn Coronado, and Don Fullerton
NBER Working Paper No. 15070
June 2009
JEL No. H22,H55

ABSTRACT

Building on the existing literature that examines the extent of redistribution in the Social Security system as a whole, this paper focuses more specifically on how Social Security affects the poor. This question is important because a Social Security program that reduces overall inequality by redistributing from high income individuals to middle income individuals may do nothing to help the poor; conversely, a program that redistributes to the poor may nonetheless be regressive according to broader measures if it also redistributes from middle to upper income households.

We have four major findings. First, as we expand the definition of income to use more comprehensive measures of well-being, we find that Social Security becomes less progressive. Indeed, when we use an "endowment" defined by potential labor earnings at the household level, rather than actual earnings at the individual level, we find that Social Security has virtually no effect on overall inequality. Second, we find that this result is driven largely by the lack of redistribution across the middle and upper part of the income distribution, so it masks some small positive net transfers to those at the bottom of the lifetime income distribution. Third, in cases where redistribution does occur, we find it is not efficiently targeted: many high income households receive positive net transfers, while many low income households pay net taxes. Finally, the redistributive effects of Social Security change over time, and these changes depend on the income concept used to classify someone as "poor".

Jeffrey R. Brown
Department of Finance
University of Illinois at Urbana-Champaign
340 Wohlers Hall, MC-706
1206 South Sixth Street
Champaign, IL 61820-9080
and NBER
brownjr@uiuc.edu

Don Fullerton
Department of Finance
University of Illinois
BIF Box#30 (MC520)
515 East Gregory Drive
Champaign, IL 61820
and NBER
dfullert@illinois.edu

Julia Lynn Coronado
Barclay's Capital
New York, NY
corolynn@gmail.com

The Social Security system is the largest government program in the U.S. today, accounting for about a quarter of all federal revenue. It involves a major tax on working individuals, and a progressive retirement benefit schedule that replaces a higher fraction of past earnings for those with low earnings. It is the most important source of income for today's elderly, comprising approximately 40 percent of all income going to individuals age 65 and over. Early architects of the Social Security program clearly intended to improve the status of the poor elderly, and were explicit that it was being designed to "prevent destitution and dependency" (Report of the Committee on Economic Security, 1935). In other words, Social Security was intended as part of the social safety net.

Has this intent been realized? This paper provides some evidence on whether Social Security helps the poor, and it shows how the answer depends on the definition of who is poor. We use a number of different definitions of "income" to classify individuals from rich to poor and to define the poverty line. Then we calculate the effects of social security on poverty. We also use a number of different definitions of "redistribution". We focus only on Old-Age Survivors Insurance (OASI), the retirement portion of Social Security. For each definition of income or redistribution, we look at changes over time. Thus, we see whether this program is doing any more or less over time to help "the poor".

Despite the progressive benefits formula, a small recent literature questions the extent to which the overall Social Security program is actually progressive. Several independent research teams have compiled evidence that Social Security's non-linear benefit formula is not sufficient to ensure overall net effects that redistribute from rich to poor.¹ Despite using different methods, different data sets, and different metrics of redistribution, these papers come to the similar conclusion that the Social Security retirement program is not as progressive as may first appear from focusing solely on the replacement rates provided by the non-linear benefit formula.

This paper starts with some of the methodological procedures of Coronado, Fullerton, and Glass (2000), but we make numerous improvements and innovations. Four contributions are most important. Relative to that paper, (i) we triple the sample size, including much better representation of the baby boom generation, (ii) we consider new methods for calculating lifetime earnings, potential earnings, social security taxes, and

¹ Some of these prior papers include Liebman (2002), Gustman and Steinmeier (2001), Cohen, Steuerle and Carasso (2001), Smith, Toder and Iams (2001), and Coronado, Fullerton and Glass (2000).

retirement benefits, (iii) we consider alternative measures of redistribution toward "the poor", and (iv) we show how each measure of redistribution changes over time, when comparing cohorts before and after the baby boom. Besides overall changes in the Gini Coefficient, we compare median net tax rates by income quintile of our sample, and the fraction of those in each quintile with negative net tax rates (positive net transfers).²

The degree of redistribution in the current system is still the subject of vigorous debate in the broader policy community, as evidenced by recent suggestions for Social Security reform. These suggestions include various forms of privatization (which might reduce the scope of redistribution by Social Security), and other reforms such as President Bush's 2005 endorsement of a shift from wage-indexing to "progressive price-indexing" (which might increase redistribution). To evaluate how any reform would affect the social safety net, however, we need a good baseline measure of the *current* system's redistribution toward the poor.

Determining the extent of income-based redistribution in the current Social Security system is a more complex exercise than it first appears, for at least four related reasons. First, as noted recently by the U.S. GAO (2004), many possible metrics – such as internal rates of return or lifetime tax rates – can be used to measure "redistribution", and each captures slightly different features of the data.

Second, the definition of "income" matters, such as whether we consider an individual based only on his or her own earnings, or based on his or her share of household earnings. Third, for any given measure of redistribution and any definition of income, the extent of redistribution may change across cohorts due to changing economic conditions, such as the increasing labor force participation rates of women. Finally, a proper accounting of the extent of redistribution must consider not just the Social Security program rules, but also a wide array of "real life" heterogeneity, such as variation across earnings levels, earnings variability, marital status, and mortality rates, just to name a few.

This paper empirically examines the extent of within-cohort income-based redistribution in the OASI retirement system. Using a micro data set on actual U.S. households, we calculate the degree of redistribution, how it varies with the measure of income employed, and how it changes from the pre-baby boom generation to the baby boom generation.

² As will be discussed below, our sample quintiles are not fully representative of the U.S. population.

In studying the retirement program, we consider several alternative definitions of income.³ We begin with a simple model of individual lifetime earnings, and then expand it to incorporate important features such as the correlation between mortality and socio-economic status, the pooling of spousal resources within a household, and “potential” earnings (which account for the fact that some individuals consume their income in the form of leisure or home production). We examine the importance of each feature of the model individually and in combination, allowing us to learn how the various features interact.

We consider three measures of redistribution: (i) the standard calculation of “effective progression,” an overall measure of the impact on inequality that is based on comparison of before-tax and after-tax Gini coefficients, (ii) comparisons of Social Security net tax rates (taxes paid minus benefits received, as a percent of income) in each sample income quintile, and (iii) the fraction of individuals in each sample income quintile that receive positive net transfers from Social Security. The first of these measures is designed to capture the *overall* redistribution of the system, which may be appropriate for those who are concerned with the overall degree of income inequality across the full income distribution. The second and third measures allow us to focus more directly on whether Social Security helps the poor. These measures are useful for those who are more concerned with the social safety net and less concerned with reducing inequality solely between households in the upper and middle parts of the lifetime distribution.

Finally, for each measure of income and each definition of redistribution, we also explore how the degree of redistribution has changed across cohorts. In particular, we compare the pre-baby-boom generation to the baby boomers. These cohorts differ along multiple dimensions, including most importantly the degree of labor force attachment of married women.

To implement this study, we use twenty-six years of data (1968-1993) from the PSID to estimate wage profiles and to construct complete lifetime earning histories for 2233 single individuals and 3780 married individuals (1890 husbands and 1890 wives). We combine simulated and actual earnings information so that each individual in the

³ The disability insurance (DI) program is an important part of the overall U.S. Social Security system, but it is conceptually distinct from the retirement system. For this reason, and due to data limitations, we follow the standard approach in this literature of focusing solely on retirement benefits. DI is highly redistributive, so including DI in the analysis would unambiguously increase the redistribution achieved by the overall system.

sample has a complete earnings profile for ages 18 – 66. The use of a core set of actual earnings observations, as opposed to relying solely on simulated or stylized earnings, allows us to capture the effects of events that may lead individuals to enter and exit the labor force, including, for example, unemployment spells. For each person, we calculate social security payroll tax in each working year and benefits received during each year of retirement, using existing Social Security rules. We thus treat each individual as if spending an entire working life under existing Social Security rules. We also incorporate information on spousal earnings and spousal benefits that are important in determining the net benefits an individual obtains from the system.

We have four major findings. First, we find that when redistribution is measured using more comprehensive concepts of income, the Social Security system exhibits less overall effect on inequality than when it is evaluated using more narrow definitions. Indeed, when evaluated using the most comprehensive measure of income (including potential earnings as well as within-household resource sharing), the Social Security retirement program exhibits virtually no overall impact, as measured by a comparison of Gini Coefficients with and without Social Security. Second, we find that this result is driven largely by the lack of redistribution between middle and higher parts of the income distribution, and under some income definitions, by regressivity of the system at higher incomes. This regressivity at the top has the effect of masking some redistribution towards the bottom of the income distribution, suggesting that the measure of redistribution matters. Third, we find that even when redistribution occurs, it is not efficiently targeted: many high income households receive net transfers, while many low income households pay net taxes. Fourth, to the extent that the effect of Social Security in reducing inequality has changed across cohorts, the change depends on the income concept employed.

We begin below by providing a brief overview of the Social Security benefit formula. While these benefit rules involve many nuances and complications, we focus on the core elements designed to make the system redistribute across income groups. In section 2, we discuss some recent literature studying Social Security's impact on overall inequality. In section 3, we describe our data sample, as well as our assumptions and methods for constructing lifetime earnings, taxes, and benefits. We discuss our measures of redistribution and lifetime income in section 4. Our primary results about Social Security's effect on poverty are reported in Section 5, while Section 6 concludes.

1. A Brief Review of the Social Security Benefit Rules

The possibility that Social Security may redistribute stems primarily from two sources. First, retirement benefits are calculated as a non-linear function of capped lifetime income, providing a higher replacement rate for those with less income. Second, the spousal and survivor benefits may accrue to individuals with very low lifetime earnings.

Under present law, the calculation of retirement benefits begins with computation of that worker's Average Indexed Monthly Earnings (AIME). The Social Security Administration keeps track of each individual's covered earnings throughout his or her lifetime. To calculate the AIME, nominal earnings for the individual in each calendar year through age 60 are multiplied by Social Security's Average Wage Index (AWI). Wages after age 60 are not indexed.⁴ The thirty-five highest years of indexed earnings (including zeros, if applicable) are then added up and divided by the number of months in 35 years (420). The resulting number is that worker's AIME.

Next, Social Security calculates the "Primary Insurance Amount" (PIA). Using the bend-points for 2006, the formula for calculating the PIA is:

$$\begin{aligned} \text{PIA} = & 0.90 \times \min[\text{AIME}, \$656] \\ & + 0.32 \times \max[0, (\min[\text{AIME}, \$3,955] - \$656)] \\ & + 0.15 \times \max[0, \text{AIME} - \$3,955] \end{aligned}$$

If an individual retires at his or her Normal Retirement Age (NRA), the basic monthly retirement benefit is the PIA.⁵ The structure of the PIA factors (0.9, 0.32 and 0.15) is such that the PIA/AIME ratio is a declining function of AIME.⁶ Thus, if two individuals are identical in all respects except for average indexed monthly earnings, the individual with the lower AIME receives a replacement rate that is at least as high as that

⁴ See <http://www.socialsecurity.gov/OACT/COLA/AWI.html> for more details of the calculation of the AWI.

⁵ For the cohort turning age 62 in the year 2005, the NRA is 66 years. In the year 2017, the NRA is scheduled to begin rising again, reaching age 67 in year 2022. In the event that a person claims benefits prior to or later than the NRA, the benefit is adjusted in a way that is approximately actuarially fair when evaluated using population life tables.

⁶ The bend point amounts in the Social Security retirement formula were \$627 and \$3,779 in 2005, but increase annually based on the average wage index (AWI). Once an individual has claimed benefits, his or her future benefit is adjusted annually to reflect changes in the Consumer Price Index (CPI). The net result of this AIME-PIA calculation and the annual indexation of the bend points is that the *initial* benefit level is indexed to wage growth, so that replacement rates remain relatively constant over time, whereas benefits after the date of claim are linked to inflation.

of the individual with the higher AIME. As a consequence, the Social Security benefit formula is often considered “progressive”.

In addition to the worker’s own retirement benefit, Social Security also provides benefits to spouses. In particular, the spouse of an insured worker is eligible to receive a benefit that is the greater of his or her own benefit (based on own past earnings) or 50 percent of the working spouse’s PIA (subject to actuarial adjustments). As we see below, these spousal benefits play an important role in the assessment of the Social Security system’s redistribution toward those with very little income.

2. The Literature on Social Security and Redistribution

Milton Friedman (1972) and Henry Aaron (1982) hypothesized that some features of Social Security may offset the redistribution of the benefit schedule when the program is evaluated on a lifetime basis. For example, annuitization is mandatory, while mortality is negatively related to income. While their focus was largely on mortality differentials, more recent work suggests that differential mortality turns out to play, at best, only a minor role in influencing overall impact on inequality (Harris and Sabelhaus, 2005).

Other features of the Social Security program, however, are shown to be more influential. Coronado, Fullerton and Glass (2000) undertake a gradual broadening of the measure of income to include lifetime income, potential income, and household income. They show that this broadening eliminates the overall impact of Social Security on inequality, as measured by the Gini Coefficients with and without Social Security transfers.

Gustman and Steinmeier (2001) also highlight the importance of accounting for family income when looking at spouse and survivor benefits. They first measure redistribution by comparing lifetime taxes and benefits across individual AIME deciles, and they then find that redistribution is roughly halved by looking at family income deciles. Using the Health and Retirement Study, which focuses on the cohort born 1931-1941, they also indicate that when families are arrayed using years in which both spouses had substantial earnings (a method of controlling for potential income), the system achieves essentially no redistribution.

In a similar vein, Liebman (2002) applies current Social Security rules to a micro-simulation model based on the 1925-1929 birth cohort from the Survey of Income and Program Participation (SIPP). While he finds that Social Security provides within-cohort

transfers equal to 13 percent of Social Security benefit payments, he provides evidence that much of that redistribution is not related to income. His research also points to the importance of spousal benefits, the role of family income, and the sensitivity of results to assumed discount rates.

Those three studies look at alternative definitions of income, and overall inequality, but they do not look more narrowly at redistribution toward the poor, and they do not look at changes over time. The next two studies look at changes over time, but do not analyze a wide variety of income measures. None of these studies focus on redistribution toward the lifetime poor.

Cohen, Steuerle and Carasso (2001) use the Modeling Income in the Near Term (MINT) model, a micro-simulation model based on the 1990-93 SIPP Survey matched with Social Security earnings records, to study the extent of redistribution in Social Security by education, race, and income. They also look at four different generations. They find that Social Security provides higher rates of return to those with lower lifetime earnings, but they point out that some specific lower-earnings groups do worse than groups with higher family income and wage rates.

Smith, Toder and Iams (2001) also use the MINT model to study the impact of the OASI program on inequality. Their primary focus is on how the redistributive effects of Social Security are changing over time, partly because of changes in tax rates and benefits, but more importantly because of changing demographics and earnings patterns in the work force. These studies, and others, are nicely summarized by the Government Accounting Office (U.S. GAO, 2004).

Here, we focus on how Social Security affects the poor on a lifetime basis, using more data, multiple definitions of income, and multiple measures of redistribution toward the poor. We also track changes in the impact of Social Security across cohorts.

3. Sample and Data Construction

At the heart of our analysis of a safety net for the lifetime poor is the calculation of the present value of each individual's lifelong stream of income, OASI taxes, and OASI benefits. Calculating these present values requires, in turn, that we have information on each worker's marital status, mortality rate, annual earnings at each age, and spouse's income and mortality. For some of our broader income measures, we also require hourly wage rates, as opposed to annual income, for both husbands and wives, so that we can

compute potential income. This section explains sources of data for these calculations.

3.1 Sample

We use the Panel Study of Income Dynamics (PSID) for all years 1968 through 1993, which provides us with up to 26 years of earnings and demographic data for a sample of the population. We include in our sample all households who appear in the PSID for at least 10 years during our sample period, and who were under age 55 in 1968. Our sample consists of over 6,000 individuals, including the over-sampling of low-income individuals.

While the core PSID sample is designed to be representative of the U.S. population, our data are not representative in three ways. First, because we are interested in studying only the OASI retirement program, rather than Disability Insurance (DI), we have removed individuals who are chronically disabled (which we define as being disabled for more than 2 years in our sample period). These individuals are likely to be covered by the DI program. They are also more likely to be non-white, single, male, and to have lower educational levels. Because these characteristics are correlated with low lifetime income, the DI program is likely to be found to redistribute toward the poor – under any definitions of redistribution and income employed in this paper. The chronically disabled comprise approximately 6-7 percent of the total sample.

A second limitation is our requirement that an individual be observed in the PSID for at least 10 years between 1968 and 1993, which requires entering the sample no later than 1984. Thus, the sample is not as ethnically diverse as the current U.S. population. For perspective, the average annual number of immigrants into the United States was over 920,000 from 1986 – 2005, which is more than twice the annual number during the 1966-1985 period.⁷ How the inclusion of immigrants would influence overall redistribution of the system is quite complex. For example, Gustman and Steinmeier (2000) show that for each year of work under the Social Security system, immigrants realize higher benefits than the U.S. born, even when their earnings are identical in all years that the immigrant has been in the U.S. This arises because of the interaction of the non-linear benefit formula and the fact that each year spent outside the U.S. is treated as a year of zero income. Accounting for such individuals in our framework would require that we have

⁷ Calculation based on Table 1 of the Department of Homeland Security Yearbook of Immigration Statistics 2005, which can be found at <http://www.uscis.gov/graphics/shared/statistics/yearbook/LPR05.htm>.

lifetime earnings not only in the U.S. but also in the immigrant's home country. It would also require that we take account of the complex "totalization agreements" that exist between the U.S. and many other nations, designed to "protect the benefits of workers who pay into the social security systems of two countries but do not earn sufficient credits to receive full benefits from one or both countries" (Barnhart, 2003).

Third, the PSID took a nationally representative sample and then added an extra national sample of low-income families. Thus the overall sample includes more than a proportional number of low-income families. Normally this problem can be overcome by use of sample weights provided by the PSID. In this case, however, we cannot use the sample weights because of the complex nature of our sample. Recall that we include in our sample any individual who appeared in 10 or more waves of the PSID, yet the PSID sample weights are specific to each year. If we arbitrarily chose one PSID year to construct weights, any individual who did not appear that year would receive no weight. Further, some individuals who *do* appear in the weighting year would not be in our sample due to the 10-year appearance requirement. In addition, the weights would not reflect the fact that we removed from the sample individuals who had a disability during their lifetime. As a consequence of not being able to construct proper weights, our sample quintiles do not represent population quintiles. The net result of the over-sampling of the poor, even after dropping the disabled (who are disproportionately poor) is that the poorest 20% in our sample represents less than 20% of the population. Nonetheless, we can still discuss different measures of income to use in the calculation of who is in that poorest group.

3.2 Lifetime Earnings Profiles

Given our sampling criterion in the PSID, we observe between 10 and 26 years of earnings data for each individual. Therefore, in order to obtain complete profiles of earnings from age 18 through age 66 for each of our sample members, we must generate out-of-sample earnings observations. We do this by estimating earnings regressions and using the estimated coefficients to generate the needed observations.

We begin by taking all observations with non-missing earnings and indexing their annual earnings by the Social Security Administration's Average Wage Index (AWI), which reflects economy-wide growth in nominal wages over time. By applying this index to all earnings in the PSID sample, we essentially examine steady state distributional

outcomes while abstracting from real economic growth. Social Security uses this index to adjust earnings, the benefit formula, and the taxable earnings cap each year, so we can arbitrarily choose any base year for our calculations. While the choice of base year may affect the level of lifetime benefits and taxes, it does not affect the ratio of benefits or taxes to lifetime earnings – which is the basis of our analysis – because both the numerator and denominator are adjusted by the same index.

Using these wage-indexed earnings profiles, we then apply a regression specification that is a modified version of the one in Bosworth, Burtless and Steuerle (2001). They impute missing earnings observations by modeling income as a step-function of age in a model with individual fixed effects. While we follow their lead by using OLS to estimate a nonlinear function of fixed effects and age, our model for estimation differs in three ways. First, as a direct control for age, we use a cubic function rather than the step-function by age interval. Second, we include additional controls for age interacted with education, gender, race, and Hispanic status. Because our specification includes individual fixed effects, the direct effect of these demographic variables is subsumed in the fixed effect. By including interactions of these variables with age, however, we allow the slope of the age-earnings profile to vary with key demographic characteristics, at the same time that the individual fixed effects allow for a person-specific intercept. We also include time varying controls for marital status and non-chronic disability status. Third, we estimate separate regressions for men, wives, and female heads of household, effectively allowing all the slope coefficients to vary across these designations.

The basic specification for each group is:

$$y_{it} = u_i + \sum_{j=1}^3 \beta_j \cdot age_{i,t}^j + \sum_{j=4}^7 \beta_j \cdot age_{i,t} \cdot Educ_{i,j-3} + \sum_{j=8}^9 \beta_j \cdot age_{i,t} \cdot Race_{i,j-7} \\ + \beta_{10} \cdot Age_{i,t} \cdot Boomer_i + \beta_{11} \cdot Married_{i,t} + \beta_{12} \cdot Disabled_{i,t} + \varepsilon_{it}$$

In this specification, y_{it} represents individual i 's income in year t , u_i is an individual fixed effect, and ε_{it} is the error term. We include age, age-squared, age-cubed, age interacted with each of four education indicators (high school, some college, college, and more than college, where the excluded category is less than high school), age interacted with two “race” indicators (non-white and Hispanic), and age interacted with a dummy variable for whether the individual is part of the baby-boom or the pre-boomer cohort. We also include time-varying controls for marital status and (non-chronic)

disability. Estimated coefficients from these regressions are in an appendix available from the authors.

Using the estimated coefficients from these regressions, we then simulate earnings for all missing years by interpolating or extrapolating the individual's age, holding constant the other characteristics such as education and race, and also including the individual fixed effect. In order to calibrate the number of zero earnings years that we expect in the out-of-sample simulation, we run Probits on a binary measure of labor force participation against age, marital status, and all other demographic characteristics. We then use these results to calculate the number of zero earnings years that we would expect in our simulated years. For heads, we achieve this by converting to zero any simulated earnings that are less than or equal to zero. For wives (female heads), we convert to zero any simulated out-of-sample earnings that are less than six (eight) forecast errors above zero.

Combining the actual observations with simulated observations yields a complete earnings profile for ages 18 to 66 for each individual in our sample. These complete earnings profiles allow us to account for entry and exit from the labor force, a factor that is important for evaluating redistribution because benefits are based on earnings histories and allow for a certain number of years to be dropped before making average wage calculations. This feature provides a major advantage over the use of stylized "average" earners that are often employed by Social Security's Office of the Chief Actuary when evaluating the distributional effects of reform.

Another advantage of using rich earnings data, relative to stylized earners, is that we have a demographically diverse sample. This diversity affects our analysis in that different demographic groups have different numbers of single and married households, different earnings patterns, and different mortality rates. These differences turn out to be an important issue in analyzing social security, as described below.

3.3 The Wage Rate and Potential Earnings

One of the measures of lifetime income that we use in our analysis is designed to account for an individual's "potential," rather than actual, earnings. We define potential income as an individual's wage rate times his or her annual endowment of potential labor hours. The wage rate is a measure of earning power that reflects, among other things, experience, talent, and education. Using an annual endowment of labor hours 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.

We regress the log of the wage rate on individual fixed effects as well as the other variables that are used as controls in the earnings regressions. As in the earnings regressions, we run separate regressions for men, wives, and female heads of household. Using the resulting fixed effects and coefficients, we fill in missing observations during the sample period and observations outside of the sample period so that each individual has a wage rate for every year of their entire economic life, from age 22 to 66.

For wives that never work in our sample, we are unable to estimate an individual fixed effect. In these cases, we assign these women a fixed effect equal to the minimum fixed effect in our sample of wives. We then use the coefficients from the regression of wives to fill in the entire profile of potential hourly wage rates. All coefficient estimates are in an appendix available from the authors.

To calculate each individual's labor endowment, we begin by assuming an annual endowment of 2000 hours (e.g., 50 weeks of labor at 40 hours per week). From this annual endowment, we subtract any hours of involuntary unemployment for that individual at each age. The age- and person-specific hours of unemployment are estimated from a Tobit specification, where the dependent variable is hours of involuntary unemployment. The dependent variables include a cubic in age, indicators for level of education, race, Hispanic status, and whether a baby boomer. As with other regressions, these Tobits are run separately for men, married women, and female heads of household, and the coefficients are used to predict unemployment for missing observations.

For each individual, we then compute the product of the estimated wage rate for each age and the individual's labor endowment at that age. The resulting number is what we term "potential earnings."

3.4 Social Security Taxes Paid

Social Security is primarily financed from the payroll tax known as FICA (Federal Insurance Contributions Act). This tax consists of three portions: Old Age and Survivors Insurance (OASI), Disability Insurance (DI), and the Medicare system's Hospitalization Insurance (HI) program. 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% of all employment in the U.S. is covered.⁸

The FICA tax is deducted from employees' pay at a rate of 7.65% of wages, but employers match that tax for a total 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 a maximum amount of taxable earnings -- the social security wage cap (\$94,200 for 2006). 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 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 tax through reduced wages. We therefore base our estimates on the combined employer and employee tax.

For a variety of reasons, we have chosen not to model the many ways in which the U.S. income tax system interacts with Social Security taxes and benefits. In effect, we look at the social security system only, without any income tax.⁹ We recognize that including such interactions would probably increase the measured progressivity of the system.

⁸ 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; some members of the clergy; 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.

⁹ Panis and Lillard (1996) point out that because the employer's portion of the payroll tax is deductible against the income tax, the net cost of the tax is lower than the full amount of the payroll tax paid. Like Panis and Lillard, we treat the entire amount of the payroll tax as the employee's cost of social security coverage. We also ignore the income taxation of Social Security benefits for those with high income in retirement, and the effects of the Earned Income Tax Credit (which is sometimes motivated on the basis that it helps to offset the burden of payroll taxes for low income workers). We take this approach for three reasons. First, we do it for comparability with other studies of Social Security's progressivity, nearly all of which have examined Social Security in isolation from the income tax system. Second, it is conceptually debatable whether the income tax affects the incidence of Social Security or Social Security affects the incidence of the income tax. An analysis of the overall tax system is clearly beyond the scope of this paper. Third, the data requirements for doing this analysis appropriately are quite large, requiring that we extrapolate income from *all* sources during retirement in order to determine the marginal tax bracket applied to Social Security benefits.

Of the total 15.3% tax, 10.6% is for Old Age and Survivors Insurance (OASI), 1.8% is for Disability Insurance (DI), and 2.9% is for Medicare (HI). The OASI portion of the tax is used to pay all retirement benefits. Because our focus is the retirement portion of the social security system, not disability insurance or hospital insurance, we ignore the DI and HI portions of the tax, as well as benefits paid from the DI and HI Trust Funds. As such, all calculations in this paper assume a 10.6% tax on earnings up to the cap.

Of course, the 10.6% payroll tax is not sufficient to finance the level of benefits scheduled under current law. As noted in the official summary of the 2006 Report of the Social Security and Medicare trustees, “Social Security can be brought into actuarial balance over the next 75 years in various ways, including an immediate increase of 16 percent of payroll tax revenues or an immediate reduction in benefits of 13 percent (or some combination of the two). To the extent that changes are delayed or phased in gradually, greater adjustments in scheduled benefits and revenues would be required. Ensuring that the system is solvent on a sustainable basis over the next 75 years and beyond would also require larger changes.”¹⁰ As we learned from the 2005 debate on this issue, however, we are quite far from having a political consensus on what mix of tax and benefit changes should be implemented to address the long-run fiscal imbalance. Thus, rather than imposing an arbitrary “reform” on the data, we calculate our measures of redistribution using currently scheduled taxes and benefits, while recognizing two limitations. First, any calculation of lifetime net tax rates using currently scheduled benefits and taxes will, on average, be “too generous” relative to what is sustainable in the long-run. Second, if changes to scheduled taxes or benefits differentially impact various points in the income distribution, this will obviously influence the extent of redistribution in the system. We leave the analysis of the distributional effects of alternative reform proposals to future work.

As noted in previous sections, our sample from the PSID includes observed and constructed earnings for each individual from ages 18 to 66. To calculate lifetime taxes, we simply multiply each year of earnings, up to the taxable maximum, by 10.6 percent. The present value of those taxes are then determined by discounting the expected tax payments, where expectations are taking with respect to survival probabilities.

3.5 Social Security Benefits

Under provisions of the Social Security Act, benefits are calculated from a non-linear formula described in section 1 of this paper. Our calculations follow the Social Security Administration's computation of AIME upon retirement. In particular, earnings prior to age 60 are indexed by the AWI for the year the individual attains age 60. Only earnings at or below the taxable cap in each year are considered. Earnings after age 60 are not indexed. A person who works from age 22 through age 66 would have a total of 45 years of earnings. Under the Act, only the highest 35 years are considered, so the ten lowest years are dropped. AIME is the simple monthly average of the indexed earnings in those 35 highest-earnings years.¹¹

As discussed in section 1, the AIME is then fed through a non-linear formula to calculate the Primary Insurance Amount (PIA). Like the cap on earnings, the bend points are adjusted annually by the proportional increase in the Average Wage Index. We calculate this PIA for 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 Primary Insurance Amount upon normal retirement age (which we assume to be age 67). A worker may still choose to retire as early as age 62, with reduced benefits.¹² 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. Our calculations below ignore these provisions for early or late retirement, as we assume workers (and their spouses) always choose the normal retirement age.

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

¹⁰ <http://www.ssa.gov/OACT/TRSUM/trsummary.html>

¹¹ 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.

¹² 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 retirement age is 67 would receive a benefit for the rest of his or her life that is reduced by 20%.

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". While we do account for survivor benefits to the spouse, we ignore non-spousal survivor benefits; in aggregate they are relatively minor.¹³

We use each individual's observed and constructed earnings profile to compute the 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.

3.6 Mortality

When calculating the expected present value of lifetime earnings, taxes or benefits, it is necessary to account for mortality probabilities of the individual. We begin by using a cohort life table for individuals entering the labor force at age 18 in the year 2006 (i.e., the 1988 birth cohort). Because mortality is correlated with variables that are themselves correlated with lifetime economic outcomes (e.g., race, education), we will also use mortality rates that are differentiated on this basis.¹⁴ Specifically, we use the mortality differentials calculated by Brown, Liebman and Pollet (2002) to adjust age and gender specific mortality rates by education and race. These mortality differentials, which were estimated using data from the National Longitudinal Mortality Survey, have been used in a number of studies on Social Security (Liebman 2002; Feldstein and Liebman 2002), annuities (Brown 2002), and wealth inequality (Kopczuk and Saez 2004).

3.7 Discount Rates

When calculating present values, we initially use a discount rate of 2% as in other prior literature. However, Caldwell, *et al* (1999) argue that the usual 2% rate is too low, because the discount rate should reflect the return that individuals could expect if they invested their contributions in real assets of comparable risk. They argue that the real safe

¹³ In 2004, a total of \$415 billion were paid from the OASI trust fund. Of that total, \$396 billion (95.4%) went to retired workers or their spouses, and only \$19 billion (4.6%) went to other survivor and miscellaneous benefits (*Annual Statistical Supplement, 2005*, Table 4A.5).

¹⁴ Mortality is related to income, since higher income individuals can afford better medical care. On the other hand, that correlation might be due to reverse causality, if a person's income falls during years of illness prior to death. In a sense, then, we use education and race as exogenous proxies for income.

return on indexed Treasury bonds is about 3.5% and that a premium should be added to reflect the fact that Social Security is not riskless. To account for this argument, we run some specifications with the discount rate increased from 2% to 4%. This change increases the net social security tax rate for everyone, because it increases the weight on earlier payments of payroll taxes relative to later receipt of benefits. Yet payroll taxes are regressive (because of the exemption of wages above the cap), and benefits are progressive (because of the formula). Thus the shift in weight from later benefits to earlier taxes is expected to reduce the overall help to those who are poor on a lifetime basis.

3.8 Equivalence Scales

When we move from analyzing individuals based on their own earnings to analyzing individuals based on their share of household earnings, we recognize that married couples typically pool their resources. When dividing household earnings across the husband and wife, we make use of the equivalence scale estimated by Citro and Michael (1995), which has also been used in numerous other papers (e.g., Scholz *et al*, 2006; Brown and Poterba 2000). This equivalence scale takes the form $n_j = A_j^{0.7}$, where A_j is the number of adults in the household.¹⁵ Thus, for a couple, we do not divide total household resources by 2 but instead divide by $2^{0.7} = 1.6245$.

4. Measures of Income and Redistribution

We focus on the reduction of poverty on a *lifetime* basis *within* a cohort. We do not analyze redistribution on an annual basis, nor inter-generational redistribution induced by pay-as-you-go financing of Social Security. In contrast to some other studies, we analyze redistribution in terms of total labor earnings, not just earnings up to the earnings cap.¹⁶

We believe it is important to capture the regressive feature of the system that the marginal

¹⁵ Relative to Coronado *et al* (2000), three of the major improvements here include (1) more data for more cohorts, (2) improved mortality adjustments, and (3) the use of equivalence scales. However, we do not adjust the equivalence scale for the presence of children in the household.

¹⁶ For example, Panis and Lillard (1996) use three hypothetical earnings groups: a “low” group at the full-time minimum wage rate, the “middle” group at the Social Security Average Earnings, and the “high” group at the wage cap. This use of these hypothetical workers implicitly ignores all earnings above the wage cap. Three or more hypothetical or arbitrary income groups are used by Myers and Schobel (1983), Hurd and Shoven (1985), Boskin, et al (1987), Steuerle and Bakija (1994), Garrett (1995), and Diamond and Gruber (1999). Actual social security records are used by Burkhauser and Warlick (1981), Hurd and Shoven (1985), Duggan, et al (1993, 1995), Liebman (2002), and Gustman and Steinmeier (2001). To estimate uncapped earnings from social security records, Fox (1982) uses information on the time of year an individual reaches the maximum. Liebman (2002) performs other imputations to assign earnings to each top-coded individual. Caldwell, et al (1999) use simulated data on earnings that are not top-coded.

tax rate drops to zero at the cap. We are able to do so because, in contrast to some data sources, PSID earnings are not top-coded.

To determine the amount of redistribution within the Social Security retirement system, one must first define what redistribution means. This requires making two key decisions. First, what redistribution metrics will we use? Second, what definition of income will we use when applying those metrics?

4.1 Measures of Redistribution

We begin by calculating the lifetime net Social Security tax rate for every individual. This lifetime net tax rate is the present value of expected OASI tax payments minus the present value of expected OASI benefits divided by the present value of the individual's lifetime income. We then use these results to calculate three different measures of redistribution. The first is an overall measure of redistribution called "effective progression" (calculated using Gini coefficients as described shortly). Second, we then calculate the overall net tax rate in each income quintile of our sample, in order to focus on just the poorest group. Third, we calculate the fraction of individuals in each sample quintile that receive positive net transfers from Social Security. We omit other measures partly to save space, and partly because of problems with some of those other measures. For example, some studies use the "internal rate of return," but this measure does not indicate how many dollars are taken from one group or given to another group. The literature includes other measures of redistribution as well, but our three measures are enough to see alternative characterizations of "redistribution".¹⁷

We first compare the overall distribution of income with and without accounting for the Social Security lifetime net tax using the metric known as "effective progression" (Musgrave and Thin 1948; Kiefer 1984). The effective progression measure is defined as:

$$EP = \frac{1 - Gini_{AT}}{1 - Gini_{BT}}$$

where $Gini_{BT}$ and $Gini_{AT}$ are the before-tax and after-tax values of the Gini coefficient, respectively, where the "tax" is the lifetime net tax rate from Social Security. As is well

¹⁷ The U.S. GAO (2004) reviews concepts and measures of redistribution and recent estimates in the literature, including calculations that compare low and high income groups with respect to: the share of total benefits received, the benefit/tax ratio, the internal rate of return, and the benefit/earnings replacement rate. Public Economics textbooks define a tax as progressive if the ratio of burden to income rises with income, and that is why we focus on the net tax as a fraction of income for each group.

known, the Gini coefficient is a measure of the inequality of a distribution, and is typically defined as the ratio of the area between the Lorenz curve of the distribution and the 45-degree line in a graph of the cumulative percentage of individuals against the cumulative percentage of income earned by those individuals. If all individuals have equal income, the Gini would be zero. Higher values of the Gini indicate higher degrees of inequality. A Gini equal to one would imply that one person had all the income.

The measure of effective progression simply compares the degree of inequality before Social Security to the degree of inequality after Social Security, holding pre-tax earnings fixed. A value of one for EP indicates that the before- and after-tax Ginis are the same, and thus that Social Security has no impact on the distribution of income. A value greater than one indicates a progressive system, while a value of less than one indicates regressivity.¹⁸

The effective progression measure is useful for understanding the *overall* impact of the Social Security system on inequality. However, one feature of Gini-based measures is that it is difficult to distinguish where in the income distribution the transfers are taking place. For example, a high degree of redistribution from the second-highest income quintile to the middle income quintile would show up as a reduction of inequality, even if the bottom of the income distribution were unaffected.

However, important policy debates are less focused on the overall degree of redistribution, and more focused on how effectively Social Security targets resources of those in the bottom of the income distribution. Put differently, caring about income inequality is not the same as caring about poverty alleviation.

To better examine the extent to which Social Security is effective at boosting the lifetime income of the poor, we also report statistics on the median lifetime net tax rate within each sample income quintile.¹⁹ To the extent that the bottom quintile of our sample

¹⁸ Keifer (1984) also reviews other indices of redistribution. Some of these use the same information as the EP measure. For example, the Pechman-Okner (1974) index is calculated as $[(Gini_{AT} - Gini_{BT})/Gini_{BT}]$. Other measures such as the Suits (1977) index are based on the tax concentration curve. It is calculated like the Gini coefficient but with the cumulative tax liability on the vertical axis plotted against cumulative income on the horizontal axis. This index is useful to analyze the incidence of pure taxes, but it cannot be used for our net social security tax rates. Since the net tax is negative for some individuals, the curve would not lie within the 1x1 box.

¹⁹ An important question, of course, is how to define who is "poor" in terms of lifetime income. This question has not been asked or discussed much in the literature. In terms of annual income, some define it on a relative basis by looking at the poorest 10% or 20%, while others define it by an absolute threshold level of income. In 2006, the U.S. Census defines as poor any single non-elderly individual with annual income less

has lifetime net tax rates from Social Security that are negative, or at least lower than those further up the income distribution, then perhaps they are net beneficiaries of the system. To provide a measure of the “efficiency” of redistribution, we also report the fraction of individuals in each sample quintile that have a negative lifetime tax rate, indicating that a net transfer is received from the Social Security system.²⁰

4.2 Definition of Income

Whichever metric we use, it is also important to determine what definition of income to use when calculating the lifetime net tax rate. The natural starting place is to consider each individual’s own lifetime earnings. Thus, the first income definition that we will use is the expected net present value of an individual’s own lifetime earnings.

Our second major step will be to replace *actual* earnings (by which we mean our combination of observed and simulated earnings) with *potential* earnings. As already noted, the use of potential earnings is meant to account for the value of leisure and home production as a way of better capturing the overall economic well-being of individuals. Consider two individuals with the same low actual lifetime earnings who receive the same net transfer from Social Security. One works only part time and spends the rest of every week productively growing her own vegetables, remodeling her own house, educating her own children, and otherwise improving her own welfare by non-market activities. The other individual has the same low market income but works long hours and therefore has none of those other production and consumption activities. One might legitimately argue that the first person is not as poor as the second, and we can only distinguish the two using our data by calculating potential additional income of the first individual. Then the same net social security transfer to that person is not going to someone as “poor” as may appear. We expect this step to reduce the measured amount of redistribution.²¹

than \$10,488. Using all of their different thresholds for families in different circumstances, they find that 12.3% of the U.S. population is poor. For lifetime income, we essentially define poverty on a relative basis, since we use the poorest quintile of our sample. Since the PSID oversamples low-income families, the bottom 20% of our sample is a similar fraction of the population as the 12.3% found by the U.S. Census.

²⁰ Given the pay-as-you go nature of Social Security, and the resulting intergenerational transfers that take place, it is not necessary for the average tax rate across the full sample to be zero.

²¹ By providing this alternative measure, we do not mean that Congress in 1935 intended to give net transfers only to those with low “potential” income. They may have been concerned only with actual income. But that does not mean that actual income is the only criterion against which we can measure redistribution today. Potential earnings represent another useful criterion, for reasons stated in the text, and so measuring the amount of redistribution on this basis can be useful for thinking about what the program really does.

The individuals most affected by this reclassification are those who spent significant time voluntarily out of the labor force, either working part-time or not at all. This would include, for example, “stay-at-home” spouses who provide valuable forms of home production without receiving formal market compensation. The logic of this approach is that these individuals have chosen to stay out of the labor force because the value of this home production is at least as high as the market wage they could receive outside the home. These individuals are now assigned higher lifetime incomes based on their earning potential. The expected result is that fewer of these individuals are classified as “poor”. Thus, Social Security is expected to have less impact on lifetime poverty.

A third major step is explicitly to account for resource sharing within households. Husbands and wives typically pool their resources, and they therefore have more similar levels of economic well-being than indicated by differences in individual earnings. The policy concern for the poor does not generally extend to the low-wage spouse of a high-wage earner. Thus, in this step, we pool the potential lifetime earnings of married individuals and divide by the equivalence scale discussed above. This change reduces income for the high-earning spouse and increases it for the low-earning spouse. Thus, the before-tax distribution of income is more equal, and net transfers by Social Security *within* a family are not considered part of “redistribution.”

5. Results: Does Social Security Redistribute?

5.1 Effective Progression

In Table 1, we report the before-tax and after-tax Gini coefficients as well as the resulting value of the Effective Progression (EP) measure. We begin by analyzing the extent of effective progression using the simplest measure of income – “actual” individual lifetime earnings (where “actual” means the combination of observed and simulated earnings constructed above). Row 1 reports the values for the case in which we apply standard mortality rates that differ only by age and gender, while row 2 reports the results incorporating additional mortality differences by education and race.

In row 1, the before-tax Gini of 0.443 drops to 0.426 after incorporating the net effects of Social Security. Applying the EP formula, this translates to an effective progression of 1.0315. Because the EP value is greater than 1, indicating a reduction in the Gini coefficient, the Social Security system can be said to reduce inequality. How much? These figures can be compared to others' results using annual income in the U.S. to

measure the effects of *all* taxes and transfers. The OECD (1995) reports a smaller Gini of 0.34 after taxes and transfers, but their income measure is top-coded (which biases the Gini downwards). Using a broader measure of annual income, Lerman and Yitzhaki (1995) calculate a Gini coefficient of 0.67 before taxes and transfers, and 0.58 afterwards. The corresponding EP measure is 1.16, but of course that includes the entire tax and transfer system. Looking only at annual individual income taxes and annual income, Keifer (1984) finds that the Gini falls from about 0.47 to 0.44 (EP=1.06). This figure would be smaller on a lifetime basis. Since our lifetime measure for Social Security alone is EP=1.03, it does appear that Social Security helps redistribute from rich to poor.

Table 1: Effective Progression

Scenario	Mortality	Lifetime Earnings Basis	Lifetime Earnings Measure	Discount Rate	Before SS Tax Gini Coefficient	After SS Tax Gini Coefficient	Effective Progression
1	Standard	Indiv.	Actual	2%	.44329	.42577	1.0315
2	Diff.	Indiv.	Actual	2%	.44616	.42903	1.0309
3	Diff.	Indiv.	Potential	2%	.30985	.29887	1.0159
4	Diff.	Household	Actual	2%	.34657	.34424	1.0036
5	Diff.	Household	Potential	2%	.27748	.27538	1.0029
6	Diff.	Household	Potential	4%	.27219	.27178	1.0006

A comparison of rows 1 and 2 indicates that incorporating mortality differences by education and race have very little effect on the results. On the one hand, this outcome might be surprising, given the substantial differences in mortality. For example, Brown (2002) shows that, conditional on reaching age 22, the remaining life expectancy difference between men and women is about 6 years, but that this difference rises to 17 years when comparing college educated white women to black males with less than a high school education. However, deeper analysis suggests that this effect is offset by two factors. First, when sorted on the basis of individual lifetime earnings, many of those in

the lowest sample quintile are married women, who in fact have better than average mortality rates. Second, we find that incorporating these mortality differentials does very little to alter each person's location in the lifetime income distribution. For example, if individuals in our sample are sorted into income quintiles on the basis of lifetime income using standard mortality rates, and then independently sorted again using differential mortality rates, we find that over 98 percent of individuals are in the same quintile under either definition. The finding that differential mortality rates do not have a first order effect is also consistent with the work of Harris and Sabelhaus (2005). We confirm that this small effect of mortality applies under other definitions of redistribution and income as well. Thus, in remaining rows, we only report results using differentiated mortality rates.

We next turn to our measure of potential earnings, which places a monetary value on non-market activities such as leisure and home production. Many individuals have high wage rates but low earnings (e.g., the college-educated stay-at-home parent), and they are now placed much further up the income distribution. Indeed, we find that only 60 percent of the individuals who are classified as being in the lowest income quintile of our sample when using actual earnings remain in the lowest income quintile when evaluated on a potential income basis. While the largest fraction of these “quintile switchers” move up just one quintile in our sample, 12.5% of the individuals who were in the lowest quintile based on individual earnings are moved to one of the top three income quintiles based on potential income.

The use of potential earnings has the effect of flattening out the earnings distribution, by raising the measure of earnings at the bottom while having little effect at the top. Thus, the Gini coefficient in row 3 is lower both before Social Security's net tax, and afterwards. The before-tax Gini falls to 0.310, while the after-tax Gini falls to 0.299. The EP measure declines to 1.016, suggesting that Social Security slightly reduces inequality, but it less clearly does so when evaluated on the basis of potential earnings.

In recognition of the fact that married couples typically pool their resources, the remaining rows of Table 1 divide family resources between husband and wife. If a low-earning wife is married to a high-earning husband, then individual level measures treat this person as a low-income individual. When we pool resources, we allow her to have access to part of her husband's resources. As noted above, we assume that each household shares its resources equally (adjusted using an equivalence scale).

Row 4 of Table 1 reports results where couples share actual household earnings. The before-tax and after-tax Gini coefficients are 0.347 and 0.344, respectively, for an EP of only 1.0036. This result indicates substantially less reduction in inequality than when using individual income measures, and it reflects the fact that much of the apparent redistribution from Social Security occurs within, rather than between, households. Indeed, the EP suggests that Social Security achieves very little net redistribution, since the presence of Social Security barely changes the Gini at all.

In row 5 of Table 1, we combine the previous two innovations, and jointly consider the concept of potential income and within-household resource sharing. These two cases do overlap, as they both increase the measured well-being of high-ability individuals who opt out of the labor force. Still, however, they are not perfectly correlated. By combining both cases, we simultaneously recognize that households share resources *and* may choose to consume some of those resources in the form of increased leisure or home production. The combined effect of these two factors is to reduce the before-tax and after-tax Gini coefficients to 0.277 and 0.275, respectively. The EP falls to only 1.0029.

We repeat this case in row 6, but using a higher discount rate of 4 percent. We find that the EP drops to only 1.0006, suggesting that the Social Security system has no effect on the overall level of income inequality when evaluated using a higher discount rate. The next subsection looks specifically at effects on the poor, rather than overall inequality.

The overall conclusion from the analysis of Gini coefficients and effective progression is that Social Security achieves virtually no reduction in inequality (or at best, very little). What little redistribution appears when evaluating the system based on individual, lifetime earnings nearly disappears when considering within-household resource sharing as well as the fact that some households choose to consume their income in the form of leisure or home production.

5.2 Does Social Security Help the Poor? Lifetime Tax Rates Across Income Groups

As noted earlier, the EP measure is designed to characterize the degree of redistribution across the entire income distribution. However, a potentially more legitimate policy concern is the extent to which Social Security does or does not help individuals at the bottom of the income distribution. The answer, of course, depends on “which income distribution?” In other words, are we interested in the bottom quintile of the distribution based on actual lifetime earnings, potential earnings, household earnings,

or a combination of potential and household earnings?

Table 2 explores the same cases as in Table 1, but for each income quintile it reports the median net lifetime tax rate. We focus on the bottom income quintile. Table 3 provides a slightly different perspective on lifetime net tax rates, reporting what fraction of individuals in each income quintile have a negative lifetime net tax rates – indicating that these households receive a net transfer from Social Security. This tells us how efficiently Social Security targets the poor. For example, if lifetime net tax rates are increasing across the quintiles, but we still find that a large fraction of individuals in the bottom income quintile have positive tax rates, while large fractions of higher income quintile individuals have negative tax rates, it would suggest that the system poorly targets those most in need.

Table 2: Median Lifetime Net Tax Rate by Income Quintile

Scenario	Mortality	Lifetime Earnings Basis	Lifetime Earnings Measure	Discount Rate	Median Net Tax Rate by Income Quintile				
					Low	Q2	Q3	Q4	High
1	Standard	Indiv.	Actual	2%	-21.94	-1.02	+3.59	+6.37	+6.78
2	Diff.	Indiv.	Actual	2%	-21.01	-1.30	+2.89	+5.09	+5.46
3	Diff.	Indiv.	Potential	2%	-2.71	-0.41	+2.26	+4.58	+4.73
4	Diff.	Household	Actual	2%	-1.32	+2.92	+3.40	+3.85	+3.61
5	Diff.	Household	Potential	2%	+0.22	+2.06	+3.08	+2.97	+2.95
6	Diff.	Household	Potential	4%	+4.10	+5.16	+5.99	+6.00	+5.51

Look first at Table 2, row 1, for actual individual lifetime income. It shows that in the lowest income quintile, the median lifetime net tax rate from Social Security is *minus* 21.9%. In other words, half of the individuals in the poorest quintile receive net transfers from Social Security that are more than 21% of lifetime earnings. In the next lowest quintile, the median net tax rate is minus 1.0%. In higher quintiles, the net tax rate is positive, ranging from 3.6% for the middle quintile to 6.8% for the top quintile. In Table 3, row 1, we see that over 86% of those in the lowest income quintile receive a net transfer,

and that this fraction is declining rapidly at higher points on the income distribution. Thus, using an individual-level measure of actual lifetime earnings suggests that the system is doing a fairly good job of targeting dollars towards the lowest income individuals. Results using differential mortality (row 2) are quite similar. This case similarly implies that social security *is* part of the social safety net.

Moving to the definition of potential income (row 3 of Table 2), the median tax rate in the bottom quintile changes from -21% to only -2.7% . For a given low-income person, “income” in the denominator is now larger, because low actual earnings are replaced by higher potential earnings. Also, the new definition has changed the composition of who is in the bottom quintile. Again, the median tax rate rises as income rises across the five quintiles. The pattern in row 3 of Table 3 suggests that the precision with which benefits are targeted to those most in need appears to diminish when using potential income.

Changing from individual to household earnings in row 4 of Table 2, it is still the case that the lowest income quintile is receiving net transfers, with a lifetime net tax rate of -1.3% . At this point, we note that two different modifications each cause a substantial reduction in the -21% subsidy rate of row 2 (with differential mortality, individual, actual earnings). First, row 3 changes only from actual to potential income, which cuts that subsidy rate to only -2.7% . Second, row 4 returns to actual income but changes only the individual to the household measure, which also cuts the -21% subsidy rate to only -1.3% . Yet when row 5 employs both of those features, household units with potential income, the change in the rate on the bottom quintile is not much more than either of those substantial effects alone (from the -21% subsidy to about zero). It seems that the combination is not equal to the sum of the parts, because these two effects interact. The reason is that either one of these modifications moves many of the same individuals out of the bottom quintile under study. For example, in scenario 2, using actual individual earnings, 60% of the individuals in the bottom sample quintile are married women. When one considers potential income, only 43% of the bottom sample quintile are married women, suggesting that many women who appear poor using actual earnings are in this group because of limited labor force participation rather than because of low wages. When one examines actual *household* earnings, only 10% of the bottom sample quintile are married women, indicating that most of the low income women are married to higher income men. When combining household and potential income concepts in scenario 5, 13% of the bottom

quintile are married women. Thus, once we control for household income, the incremental effect of also considering potential earnings is smaller.²²

Next, we look at how the relative treatment of the highest two income groups changes from row 3 to row 4 of Table 2, when evaluated on a household basis: the median tax rate in the 4th decile is somewhat higher than the median tax rate in the top decile. In other words, the system is regressive at the top end of the income distribution. This result could occur if, for example, the top quintile includes a large fraction of single-earner married couples (who, due to spousal benefit rules, get a higher “return” on their contributions) while the fourth quintile contains more dual-earner couples (who tend to receive a lower return on their contributions). It can also occur because of the regressive nature of the earnings cap, which limits the exposure of high earners to the payroll tax. Table 3 also shows a further decline in the precision with which net transfers are targeted to those in the bottom of the income distribution.

When using potential household income (row 5 of Table 2), we find that the median lifetime net tax rate on the poorest quintile is now just slightly positive at 0.22%. The net tax rate is larger at higher income quintiles, but interestingly, it does not vary by much across the upper quintiles and is again slightly regressive. An examination of row 5 in Table 3 shows that less than half of those in the bottom income quintile receive positive transfers from Social Security, hardly a "safety net", while 12-16% of those in the upper three income quintiles do receive positive transfers. Thus, while Social Security may, on average, transfer some resources to the lowest income quintile, the program is not well targeted. It both (a) fails to redistribute towards a large fraction of those most in need, and (b) unnecessarily transfers resources to large numbers of high income individuals.

As expected, the use of a higher discount rate in row 6 of Table 2 raises net tax rates for everyone (because the benefits are much farther into the future than are the tax payments, and thus they are discounted more heavily). In Table 2, all five quintiles have median net tax rates between 4% and 6%. Because the higher discount rate raises net tax rates across the board, we see in Table 3 that the fraction of individuals receiving net transfers drops across the board.

²² Similar patterns also occur when the sample is examined using other demographic characteristics, such as race or educational attainment.

Table 3: Fraction of Individuals Receiving Net Transfers from Social Security

Scenario	Mortality	Lifetime Earnings Basis	Lifetime Earnings Measure	Discount Rate	% of Members in Each Quintile with Social Security Net Tax Rate < 0				
					Low	Q2	Q3	Q4	High
1	Standard	Indiv.	Actual	.02	86.4	57.5	14.6	1.5	0.2
2	Diff.	Indiv.	Actual	.02	79.7	58.1	24.1	4.5	4.4
3	Diff.	Indiv.	Potential	.02	65.5	53.1	31.2	16.1	9.6
4	Diff.	Household	Actual	.02	56.9	27.0	14.1	6.9	14.4
5	Diff.	Household	Potential	.02	47.7	30.0	16.3	11.8	13.5
6	Diff.	Household	Potential	.04	15.4	3.8	1.3	0.8	0.7

A comparison of the net tax rate results with the effective progression results allows us to develop a better understanding of the underlying dynamics. The key feature is that, even when the measure of effective progression indicates very little overall impact on inequality, individuals in the lowest income quintile still have, on average, a lower lifetime net tax rate from Social Security than do individuals in higher income quintiles. In the upper half of the income distribution, however, the system exhibits little redistribution, and indeed some evidence of regressivity. Of course, even when focusing solely on the net tax rates for the lowest income quintile, it remains the case that the system appears less redistributive when using broader income measures.

Together, these results suggest three main conclusions. First, once we account for within household resource sharing as well as voluntary time allocation to home production and leisure, the impact of Social Security on overall inequality nearly disappears. Second, these results suggest that while Social Security is not particularly good at flattening out the overall income distribution, it nonetheless is at least mildly successful at transferring resources, on average, to the lifetime poor. Third, even when Social Security is successful at targeting resources to the lowest quintile on average, many low-income households still

pay net taxes and many high income households still receive net transfers. As such, the transfers that are made by Social Security are not always well-targeted to those most in need. Whether Social Security's tax and benefit structure could be reformed to target resources to the lifetime poor more efficiently is an interesting question for future research.

5.3 Is the Degree of Redistribution Changing?

The previous section demonstrates that a shift from actual to potential earnings, or a shift from individual to household resources, strongly influences the degree of measured redistribution in the Social Security system. Given that these effects are driven, at least in large part, by the labor force participation patterns of workers (and in particular, spouses of high earners), an interesting question is whether this pattern of results changes along with labor force participation patterns over time. For example, it is well known that labor force participation rates of women have increased dramatically over the last 50 years. Thus, to the extent that we see fewer one-earner couples and more two-earner couples, the extent of measured redistribution in the system might plausibly be expected to change.

To address this question, we split our PSID sample into two sub-samples based on birth years. The first sub-sample is our "pre-Boomer" cohort, namely, those who were born prior to 1946. The second is our "Baby Boomer" cohort, namely those born in 1946 or after. By repeating the analysis of section 5 on these two sub-samples separately, we can learn to what extent the different labor market attachments of these two groups influence the degree of redistribution.

In Table 4, we report Effective Progression results for the full sample, the pre-boomer sample, and the baby boomer sample under each of our six scenarios. Using the traditional concepts of individual lifetime income in the first three rows, it appears that Social Security is becoming less redistributive over time, because the EP measure is larger for the Pre-Boomer generation than for the Baby Boom generation. When income is evaluated on a household basis, however, that conclusion is reversed. Indeed, in the pre-Boomer sample, the system actually appears to be *regressive* overall, with an EP below 1 for two of the cases. Later, for the baby boomer sample, the EP is positive (but very small). A similar story, that Social Security is becoming slightly less redistributive when using individual income measures and slightly more redistributive when using household measures, is seen by looking either at the lifetime tax rates for the five income quintiles or

at the fraction of negative lifetime tax rates in each quintile.²³ The perceived decline in overall redistribution based on individual income, as one compares pre-boomers and baby boomers, is reversed when examining the data using broader income concepts.

Table 4: Effective Progression for Pre-Boomers Compared to Baby Boomers

Scenario	Mortality	Lifetime Earnings Basis	Lifetime Earnings Measure	Discount Rate	Effective Progression Full Sample	Effective Progression Pre-Boomer Sample	Effective Progression Baby-Boomer Sample
1	Standard	Indiv.	Actual	.02	1.0315	1.0371	1.0286
2	Diff.	Indiv.	Actual	.02	1.0309	1.0373	1.0276
3	Diff.	Indiv.	Potential	.02	1.0159	1.0181	1.0146
4	Diff.	Household	Actual	.02	1.0036	0.9991	1.0070
5	Diff.	Household	Potential	.02	1.0029	1.0003	1.0047
6	Diff.	Household	Potential	.04	1.0006	0.9974	1.0025

Overall, however, the examination of redistribution within either generation tells a very similar story to that of the overall sample. When using the most inclusive concept of income that accounts for the earnings potential of both head and spouse, the Social Security system does not appear to reduce inequality in any meaningful way.

6. Conclusion

In this paper we seek to measure the extent to which the current social security system supports the social safety net, that is, the extent to which it redistributes resources to the poor. To do so, we build a model that incorporates all the information needed to categorize individuals by lifetime resources and to calculate their taxes paid and benefits received from the system. We have several findings.

First, we find that using comprehensive concepts of income implies that the Social

²³In the interest of space, these tables are not included here, but are available from the authors upon request.

Security system exhibits less overall redistribution than when it is evaluated using more narrow definitions of income. For example, when evaluated using potential labor earnings at the household level (rather than actual individual earnings), the Social Security retirement program exhibits virtually no overall impact on inequality (as measured by the change in the Gini coefficient). Second, we look at each quintile of the lifetime income distribution. We find that the lack of impact on overall inequality is largely driven by the lack of impact across the middle and upper part of the income distribution, whereas most of those in the bottom income quintile may, in fact, still get net benefits from the program. Third, even when redistribution does occur, we find that it is not efficiently targeted, with many high income households receiving net transfers, while many low income households pay net taxes. Finally, we show that the impact of Social Security on inequality differs across cohorts, but the direction and extent of those changes depends on the income concept employed.

This research suggests several areas for future work. First, this analysis explicitly ignores behavioral responses to the Social Security system, including changes in labor supply and/or savings behavior that might influence how we think about the system's impact on inequality. Second, this analysis is purely financial, and it thus misses important insurance aspects of the Social Security program, including its provision of earnings insurance, longevity insurance, and disability insurance. Analysis of the insurance value of Social Security would require a sophisticated dynamic programming model that embeds this analysis in a utility framework. Finally, given the poor long-term fiscal outlook of the Social Security program, this framework could be used to explore the distributional implications of alternative reform options.

References

- Aaron, Henry J. 1982. *Economic Effects of Social Security*. Washington DC: The Brookings Institution.
- Barnhart, Jo Anne B. "Testimony by Commissioner Jo Anne B. Barnhart, Hearing on International Social Security Agreements, Subcommittee on Immigration, Border Security, and Claims." September 11, 2003. Available at http://www.ssa.gov/legislation/testimony_091103.html.
- Bosworth, Barry, Gary Burtless and Eugene Steuerle. 2001. "Lifetime Earnings Patterns, the Distribution of Future Social Security Benefits, and the Impact of Pension Reform." *Social Security Bulletin*, Vo. 63, No. 4 (September).

- Brown, Jeffrey. 2002. "Differential Mortality and the Value of Individual Account Retirement Annuities." In M. Feldstein and J. Liebman, The Distributional Effects of Social Security Reform. University of Chicago Press: Chicago, IL.
- Brown, Jeffrey R., Jeffrey B. Liebman, and Joshua Pollet. 2002. "Estimating Life Tables that Reflect Socioeconomic Differences in Mortality," in M. Feldstein and J. Liebman, The Distributional Effects of Social Security Reform, University of Chicago Press: Chicago, IL, 2002. pp. 447 – 457.
- Brown, Jeffrey R. and James M. Poterba. 2000. "Joint Life Annuities and the Demand for Annuities by Married Couples," *The Journal of Risk and Insurance*, Vol. 67, No. 4, December, pp. 527-553.
- 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, Jagdeesh Gokhale, Thomas Johnson, and Laurence J. Kotlikoff. 1999. "Social Security's Treatment of Postwar Americans." *Tax Policy and the Economy* 13, Cambridge, MA: MIT Press.
- Cohen, Lee, C. Eugene Steuerle, and Adam Carasso. 2001. "Social Security Redistribution by Education, Race, and Income: How Much and Why." Tax Policy Center Research Report.
- Citro, Constance F., and Robert T. Michael. 1995. Measuring Poverty: A New Approach. National Academy Press, Washington DC.
- Coronado, Julia Lynn, Don Fullerton, and Thomas Glass. 2000. "The Redistribution of Social Security." NBER Working Paper #7520.
- Diamond, Peter and Jonathan Gruber (1999). "Social Security and Retirement in the U.S." in J. Gruber and D. Wise, eds., *Social Security and Retirement Around the World*. Chicago: University of Chicago Press.
- Duggan, James E., Robert Gillingham and John S. Greenlees. 1993. "Returns Paid to Early Social Security Cohorts." *Contemporary Policy Issues* 11 (October): 1-13.
- Duggan, James E., Robert Gillingham and John S. Greenlees. 1995. "Progressive Returns to Social Security? An Answer from Social Security Records." Washington, DC: Treasury Department, Office of Economic Policy, Research Paper #9501.
- Feldstein, Martin and Jeffrey B. Liebman. 2002. "The Distributional Aspects of an Investment-Based Social Security System." , in M. Feldstein and J. Liebman, The Distributional Effects of Social Security Reform, University of Chicago Press: Chicago, IL.
- Friedman, Milton. 1972. "Second Lecture." in W. Cohen and M. Friedman, *Social*

- Security: Universal or Selective?* Washington, DC: American Enterprise Institute.
- Garrett, Daniel M. 1995. "The Effects of Differential Mortality Rates on the Redistribution of Social Security." *Economic Inquiry* 7 (July): 457-475.
- Gustman, Alan L. and Thomas L. Steinmeier. 2001. "How Effective is Redistribution Under the Social Security Benefit Formula?" *The Journal of Public Economics* 82.
- Gustman, Alan L. and Thomas L. Steinmeier. 2000. "Social Security Benefits of Immigrants and U.S. Born." In George Borjas, editor, *Issues in The Economics Immigration*, University of Chicago Press, pp. 309-350.
- Fox, Alan. 1982. "Earnings Replacement Rates and Total Income: Findings from the Retirement History Study." *Social-Security-Bulletin* 45 (October): 3-23.
- Hamermesh, Daniel S. and Albert Rees. 1993. *The Economics of Work and Pay*. New York: Harper Collins College Publishers.
- Harris, Amy and John Sabelhaus. 2005. "How Does Differential Mortality Affect Social Security Finances and Redistribution?" Congressional Budget Office Working Paper 2005-5 (May).
- Hurd, Michael D. and John B. Shoven. 1985. "The Distributional Impact of Social Security." in D. Wise, ed., *Pensions, Labor, and Individual Choice*, Chicago: University of Chicago Press for NBER.
- Kiefer, Donald W. 1984. "Distributional Tax Redistribution Indexes." *National Tax Journal* 37 (December): 497-513.
- Kopczuk, Wojciech and Emmanuel Saez. 2004. "Top Wealth Shares in the United States, 1916-2000: Evidence from Estate Tax Returns." NBER working paper 10399, March.
- Lerman, Robert I. and Shlomo Yitzhaki. 1995. "Changing Ranks and Inequality Impacts of Taxes and Transfers." *National Tax Journal* 48 (March): 45-59.
- Liebman, Jeffrey B. 2002. "Redistribution in the Current U.S. Social Security System." M. Feldstein and J. Liebman, eds., *The Distributional Effects of Social Security Reform*. Chicago: The University of Chicago Press for NBER.
- Musgrave, R.A. and T. Thin. 1948. "Income Tax Progression 1929-1948." *Journal of Political Economy* 56 (December): 498-514.
- 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.
- Organisation for Economic Cooperation and Development. 1995. *Income Distribution in OECD Countries*. Social Policy Studies No. 18.
- Panis, Constantijn W.A. and Lee A. Lillard. 1996. "Socioeconomic Differentials in the Returns to Social Security." Santa Monica, CA: RAND Working Paper #96-05.

- Pechman, Joseph A. and Benjamin A. Okner. 1974. *Who Bears the Tax Burden?* Washington, DC: The Brookings Institution.
- Report of the Committee on Economic Security. 1935. Text of report is available at <http://www.ssa.gov/history/reports/ces/ces5.html>
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving Optimally for Retirement?" *Journal of Political Economy*. Vol. 114, No. 4 (August). Pp. 607-643.
- Smith, Karen E., Eric Toder, and Howard Iams. 2001. "Lifetime Distributional Effects of Social Security Retirement Benefits." Tax Policy Center Research Report.
- 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.
- Suits, Daniel B. 1977. "Measurement of Tax Redistribution." *American Economic Review* 67 (September): 747-752.
- U.S. General Accounting Office. 2004. "Social Security: Distribution of Benefits and Taxes Relative to Earnings Level." GAO-04-747.
- U.S. Social Security Administration. 2005. *Annual Statistical Supplement, 2005*. Social Security Bulletin. Washington DC: U.S. Government Printing Office.