

The Progressivity of Social Security

Jeffrey R. Brown
University of Illinois at Urbana-Champaign
and NBER

Julia Lynn Coronado
Barclay's Global Investors

Don Fullerton
University of Texas at Austin
and NBER

December 15, 2006

Abstract: This paper provides new evidence on the progressivity of the Social Security retirement program. Using the PSID, we explore how progressivity differs depending on the definition of income, how it is affected by the measure of progressivity used, and how progressivity may change over time due to differences in the economic behavior of successive cohorts. We have four major findings. First, we find that when progressivity is measured using more comprehensive concepts of income, the Social Security system exhibits less overall progressivity than when it is evaluated using more narrow definitions of income. Indeed, when evaluated using potential labor earnings at the household level (rather than actual earnings at the individual level), the Social Security retirement program exhibits virtually no overall progressivity as measured by the change in the Gini coefficient. Second, we find that this result is largely driven by the lack of progressivity (and in some cases, the presence of regressivity) in the middle and upper part of the income distribution, which masks the presence of positive, if small, net transfers to the bottom income quintile. Third, we find that even when there is redistribution occurring, 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 extent to which progressivity differs across cohorts depends on the income concept used.

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 NBER. This paper is a substantially revised and extended version of Coronado, Fullerton and Glass "The Progressivity of Social Security" NBER WP #7520.

Key words: Social Security, progressivity, redistribution

1. Introduction

The Social Security system is the largest government program in the United States today, accounting for approximately one-quarter of all federal revenue. It is also the most important source of income for today's elderly, comprising approximately 40 percent of all income going to individuals age 65 and over. This massive program is generally thought to be an important element of the social safety net in the U.S. Early architects of the Social Security program clearly intended the program to improve the status of the poor elderly, and were explicit that it was the program was being designed to "prevent destitution and dependency" (Report of the Committee on Economic Security, 1935). In addition to helping to target resources towards the poorest seniors, the program is viewed as being "progressive" with respect to earnings more generally, due to the fact that, all else equal, its non-linear benefit formula provides a higher replacement rate to individuals with lower lifetime earnings.

In recent years, however, a small literature has emerged which has begun to question the extent to which the Old-Age Survivors Insurance (OASI) program, the retirement portion of the program, is successful at redistributing income in a progressive manner. In the past few years, several independent research teams, including Liebman (2002), Gustman and Steinmeier (2001), Cohen, Steuerle and Carasso (2001), Smith, Toder and Iams (2001), and the direct predecessor of this paper (Coronado, Fullerton and Glass 2000)¹, have begun compiling evidence that Social Security's non-linear benefit formula is not sufficient to ensure overall progressivity. Despite using different methods, different data sets, and different metrics of progressivity, these papers came to the similar conclusion that the Social Security retirement program is not as progressive as it may first appear when one focuses solely on the replacement rates provided by the non-

¹ Relative to the original Coronado, Fullerton and Glass (2000) working paper, the current paper is revised and extended in numerous ways, including, but not limited to: (i) a tripling of the sample size, including much better representation of the baby boom generation, (ii) new methods for calculating lifetime earnings, potential earnings, benefits and taxes, (iii) an analysis of median net tax rates by income quintiles, (iv) new analysis on how

linear benefit formula.

The degree of progressivity in the current system is still the subject of vigorous debate in the broader policy community, as evidenced by the 2005 national debate about Social Security reform. Much of the policy analysis, as well as political rhetoric, surrounding President Bush's endorsement of a shift from wage-indexing to "progressive price-indexing" centered around how individuals at different points in the income distribution would be affected by reform. A necessary precondition for evaluating how any reform will affect overall progressivity, however, is to have a good baseline measure of how progressive the *current* system is.

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 US GAO (2004), there are many possible metrics – such as internal rates of return, lifetime tax rates, etc. – that one can use to measure progressivity, each of which 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 whether we consider spousal earnings as well. Third, for any given measure and any definition of income, the extent of progressivity 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 in 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 progressivity, how it varies with the measure of income used, and how it varies

progressivity varies across the pre-baby boom and the baby boom cohorts, as well as numerous other improvements and innovations.

across the pre-baby boom generation and the baby boom generation.

In studying the retirement program,² we consider several alternative definitions of income. We begin with a simple model of individual lifetime earnings, and then expand from there to incorporate important features such as the correlation between mortality and socioeconomic status, “potential” earnings (which accounts for the fact that some high income individuals may choose to consume their income in the form of non-labor market activities), and the pooling of spousal resources within a household. 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 progressivity: (i) a measure of “effective progression,” which is based on the comparison of before- and after-tax Gini coefficients, (ii) comparisons of Social Security net tax rates (taxes paid minus benefits received as a percent of income) across income quintiles, and (iii) the fraction of individuals in each income quintile that receive positive net transfers from Social Security. The first of these measures is designed to capture the degree of *overall* progressivity of the system, which may be an appropriate measure for those who are most concerned with the overall degree of income inequality across the full income distribution. The second and third measures allow one to focus more directly on whether Social Security helps the lowest income individuals. These measures are useful for those who are more concerned with transferring resources to the lifetime poor, and less concerned with redistribution between households in the upper and middle parts of the lifetime resources distribution.

Finally, for each measure of income and each definition of progressivity, we also explore how the degree of progressivity differs across cohorts. In particular, we compare the pre-baby

² We note at the outset that the disability insurance (DI) program is an important part of the U.S. Social Security system. DI is, however, a conceptually distinct insurance program from the retirement system. For this reason, and due to data limitations, we follow the standard approach in the literature of focusing solely on the retirement benefits from the system. Because DI is highly redistributive, including DI in the analysis would unambiguously increase the progressivity of the overall Social Security system.

boom generation to the baby boom generation, which differ along multiple dimensions, the most important of which is 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 3780 married individuals (1890 husbands and 1890 wives) and 2233 single individuals. We combine simulated and actual earnings information so that each individual in the 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 paid during each year of retirement using existing Social Security rules, thus treating each individual as if they had spent their 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 progressivity is measured using comprehensive concepts of income, the Social Security system exhibits less overall progressivity than when it is evaluated using more narrow definitions of income. Indeed, when evaluated using the most comprehensive measure of income (which accounts for both potential earnings as well as within-household resource sharing), the Social Security retirement program exhibits virtually no overall progressivity as measured by a comparison of Gini coefficients before and after accounting for the presence of Social Security. Second, we find that this result is driven largely by the lack of progressive redistribution middle and higher income quintiles, and under some income definitions, by the actual regressivity of the system at higher incomes. This regressivity at the top has the effect of masking some progressive redistribution towards the

bottom quintile, suggesting that the measure of progressivity matters. Third, we find that even when there is redistribution occurring, it is not efficiently targeted, with many high income households receiving net transfers, while many low income households pay net taxes. Fourth, we show that the extent to which progressivity differs across cohorts depends on the income concept used.

This paper proceeds as follows. In section 2, we begin by providing a very brief overview of how the Social Security benefit formula operates. While there are many nuances and complications to the benefit rules, we focus on the core elements that are designed to make the system redistributive along income lines. In section 3, we provide a brief overview of the recent literature on Social Security progressivity. In section 4, we provide an overview of our data sample, as well as our assumptions and methods for constructing lifetime earnings, taxes and benefits. We discuss our measures of progressivity and lifetime income in section 5. Our primary results about the progressivity of Social Security are reported in Section 6. Section 7 concludes.

2. A Brief Review of the Social Security Benefit Rules

The possibility that Social Security may be progressive stems from the fact that benefits are calculated as a non-linear function of (capped) lifetime income. Under present law, the calculation of a worker's retiree benefit begins with computation of the worker's Average Indexed Monthly Earnings (AIME). The Social Security Administration keeps track of each individual's covered earnings throughout one's 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 (wages after age 60 are not indexed).³ The thirty-five highest

³ More details of the calculation of the AWI are available at <http://www.socialsecurity.gov/OACT/COLA/AWI.html>

years of indexed earnings (including zeros, if applicable) are then added up and divided by 420 (the number of months in 35 years). The resulting number is that worker's AIME.

Next, Social Security calculates the "Primary Insurance Amount," or PIA. In 2006, the formula for calculating the PIA was:

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

If an individual retires at their Normal Retirement Age (NRA), their basic monthly retirement benefit is equal to the PIA.⁴ 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.⁵ Thus, if two individuals are identical in all respects except for their average indexed monthly earnings, the individual with the lower AIME will receive a replacement rate that is greater than (or equal to) the individual with the higher AIME. As a result, the Social Security benefit formula is often considered "progressive," and indeed this characterization is correct so long as one is holding all else equal.

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 the their own benefit (based on own past earnings) or 50 percent of the working spouse's PIA (subject to actuarial adjustments). As we will see below, these spousal benefits will play an important role in the assessment of the Social Security system's progressivity.

⁴ 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 one claims benefits prior to or later than the NRA, their benefit is actuarially adjusted, and this adjustment is approximately actuarially fair when evaluated using population life tables.

⁵ The bend point amounts in the Social Security retirement formula, equal to \$627 and \$3,779 in 2005, increase annually based on average wage growth (AWI). Once an individual has claimed benefits, their 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 claiming are linked to inflation.

3. The Literature on Social Security and Redistribution

Milton Friedman (1972) and Henry Aaron (1982) hypothesized that some features of Social Security, such as mandatory annuitization in the presence of an income-mortality correlation, may offset the progressivity of the benefit schedule when the program is evaluated on a lifetime basis. While their focus was, in large part, on the role of mortality differentials, more recent work suggests that differential mortality turns out to play, at best, only a minor role in influencing overall progressivity (Harris & Sabelhaus 2005).

Other features of the Social Security program, however, have been shown to be more influential. In an earlier incarnation of this project, Coronado, Fullerton and Glass (2000) showed that a gradual broadening of the measure of income to include lifetime, potential, household income eliminated the overall progressivity of Social Security as measured by the change in Gini coefficients before and after considering the role of Social Security.

Gustman & Steinmeier (2001) also highlight the importance of accounting for spouse and survivor benefits, arguing that redistribution (as measured by a comparison of lifetime taxes and benefits across AIME deciles) is roughly halved when these are taken into account and redistribution is measured among families. Using the Health and Retirement Study, which focuses on the cohort born 1931-1941, they also indicate that when families are arrayed by years in which both spouses had substantial earnings (a method of controlling for potential income), there was essentially no redistribution remaining in the system.

In a similar vein, Liebman (2002), using a micro-simulation model based on the 1925-1929 birth cohort from the Survey of Income and Program Participation (SIPP) and current Social Security rules, provides evidence that while Social Security provides within-cohort transfers of 13 percent of Social Security benefit payments, much of the 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.

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 find that Social Security does provide higher rates of return to those with lower lifetime earnings, although 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 redistribution in the OASI program. 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 U.S. Government Accounting Office (GAO, 2004).

4. Sample and Data Construction

At the heart of our analysis of progressivity 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 annual earnings of the worker at each age, mortality rates, marital status, and the 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 may compute a measure of potential income. This section explains the sources of data for these calculations.

4.1 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 6000 individuals, including the over-sampling of low-income individuals.

While the PSID sample is designed to be representative of the U.S. population, there are two ways in which this data is not representative. First, because we are interested in studying only the OASI retirement program, rather than DI, we have removed individuals who are chronically disabled (which we define as being disabled for more than 2 years in our sample period), because these individuals are likely to be covered by the DI program. These individuals are 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 progressive using any combination of metric and income employed in the paper. The chronically disabled comprise approximately 6-7 percent of the total sample.

A second data limitation is that because an individual must be observed in the PSID for at least 10 years between 1968 and 1993, thus requiring that they enter the sample no later than 1984, the sample is not as ethnically diverse as the current U.S. population. For perspective, from 1986 – 2005, the average number of annual immigrants into the United States was over 920,000, versus less than half this number annual during the 1966-1985 period.⁶ How the inclusion of immigrants would influence overall progressivity 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 years an immigrant spends outside the US are treated as zero years of income. A proper accounting of such individuals in our framework would require not only that we have lifetime earnings in the U.S., but also lifetime earnings in the immigrant's home country. It would also require that we take account of the

⁶ 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>.

complex “totalization agreements” that exist between the U.S. and many other nations, which are 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).

4.2 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 can, in essence, examine the steady state distributional outcomes while abstracting from real economic growth. Because Social Security uses this index to adjust earnings, the benefit formula, and the taxable earnings cap each year, we can arbitrarily choose any base year for our calculations. While the choice of the year will affect the level of lifetime benefits and taxes, it will not effect 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 that approached used by Bosworth, Burtless and Steuerle (2001) who 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 of estimating an OLS model as a function of fixed effects and age, our estimation method differs in three ways. First, as a direct

control for age, we use a cubic function as opposed to the step-function by age interval. Second, we include additional controls for education, race, Hispanic status and gender, each interacted with age. Because our specification includes individual fixed effect, 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 and age cubed, age interacted with each of four education indicators (high school, some college, college, college+, with less than high school as the excluded category), 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 vs. the pre-boomer cohort. We also include time varying controls for marital status and (non-chronic) disability. More details on these regressions, including how we dealt with earnings skewness, can be found in Appendix A.

Using the estimated coefficients from these regressions, we then simulate earnings for all missing years by interpolating or extrapolating the individual's age, holding other characteristics (education, race, etc.) constant, 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

ran Probits on a binary measure of labor force participation against age, marital status, and all other demographics. We then used these results to estimate the number of zero earnings years that we would expect in our simulated earnings. For heads, we achieved this by converting to zero any simulated earnings that were less than or equal to zero. For wives (female heads), we converted to zero any simulated out-of-sample earnings that were less than six (eight) forecast errors above zero. More details on this process are provided in Appendix A.

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 progressivity because benefits are based on earnings histories and allow for a certain number of years to be dropped before making average wage calculations. This is a major advantage over the use of stylized “average” earners that are often used 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.

4.3 Wages and Potential Earnings

One of the measures of lifetime income that we will 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 their 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 wage on individual fixed effects as well as the other variables that were 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 wages.

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."

4.4 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.⁷

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.⁸

⁷ 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.

⁸ 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, and for comparability with other studies, we treat the entire amount of the payroll tax as the employee's cost of social security coverage. In effect, we look at the social security system only, without any income tax. The combined

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.”⁹ 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 progressivity 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.

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.

As noted in the 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.

4.5 Social Security Benefits

Under provisions of the Social Security Act, benefits are calculated from a non-linear formula that was described in section 2 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 will be dropped. AIME is the simple monthly average of the indexed earnings in those 35 highest-earnings years.¹⁰

As discussed in section 2, 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 (which we assume to be age 67). A worker may still choose to retire as early as age

⁹ <http://www.ssa.gov/OACT/TRSUM/trsummary.html>

¹⁰ 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.

62, with reduced benefits.¹¹ 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 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.¹²

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.

4.6 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

¹¹ 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%.

¹² In 2004, a total of \$415 billion were paid from the OASI trust fund. Of that total, \$396 billion (95.4%) went to

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. 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).

4.7 Discount Rates

When calculating present values, we will initially use a discount rate of 2%. However, Caldwell, et al (1999) argue that the 2% rate used in much of prior literature 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 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 will show 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 overall progressivity.

4.8 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

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*).

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 & 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.¹³ As such, rather than dividing total household resources by 2, we divide by $2^{0.7} = 1.6245$.

5. Measures of Income and Redistribution

The focus of this study is an examination of redistribution on a *lifetime* basis *within* a cohort. We are not focusing on the extent of redistribution on an annual basis, nor are we examining the extent of intergenerational redistribution that the pay-as-you-go financing structure of Social Security induces. In contrast to some previous studies, we are also going to think of redistribution in terms of total labor earnings, not just those subject to the earnings cap.¹⁴ We believe it is important to capture the regressive feature of the system that the marginal 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 how progressive the Social Security retirement system is, one must first define what progressivity 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?

¹³ We do not adjust the equivalence scale for the presence of children in the household.

¹⁴ 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.

5.1 Measures of Progressivity

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 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 progressive 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 known, the Gini 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 the cumulative percentage of people 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.¹⁵

¹⁵ Kiefer (1984) also reviews other indices of progressivity. 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

The effective progression measure is useful for understanding the *overall* progressivity of the Social Security system. 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, if there were a high degree of redistribution from the 2nd highest income quintile to the middle income quintile, this would show up as an increase in overall progressivity, even if the bottom two quintiles were unchanged.

However, some policy debates are less focused on the overall degree of redistribution, and more focused on how effective Social Security is at targeting resources at those in the bottom of the income distribution. Put differently, caring about income inequality is not exactly the same thing 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 rates by income quintile. To the extent that the bottom quintile has lifetime net tax rates from Social Security that are negative, or at least lower than those for higher income quintiles, this suggests that they are net beneficiaries of the system. To provide a sense for how “efficient” any redistribution is, we also report what fraction of individuals in each quintile have a negative lifetime tax rate, indicating a net transfer from Social Security.¹⁶

5.2 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 next major step will be to replace *actual* earnings (by which we mean our combined

security tax rates. Since the net tax is negative for some individuals, the curve would not lie within the 1×1 box.

¹⁶ Given the pay-as-you go nature of Social Security, and the resulting intergenerational transfers that take place, it is

actual 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.

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 people are now assigned higher lifetime incomes based on their earning potential. The result should be that the entire distribution of before-tax lifetime income is now more evenly distributed, and Social Security would thus be expected to have a less-progressive effect.

A third major step will be to explicitly 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, we now pool the potential lifetime earnings of married individuals and divide by the equivalence scale noted 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.”

6. Results: Is Social Security Progressive?

6.1 Effective Progression

In table 1, we report the before 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.

The before-tax Gini in row 1 of 0.443 drops to 0.426 once one incorporates the Social Security net tax rate. 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 be progressive. How progressive? These figures can be compared to others 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. Looking only at individual income taxes, Keifer (1984) finds that the Gini falls from about 0.47 to 0.44 (EP=1.06). Thus, it appears that Social Security is progressive on a lifetime basis, although the extent of redistribution is likely less than that observed in the income tax system.

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, such a result might be surprising, given that the substantial differences in mortality. For example, Brown (2002) shows that, conditional on reaching age 22, the remaining life expectancy between men and women is about 6 years, but that this difference rises to 17 years when one compares a college educated white woman to a black male 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 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 one's location in the lifetime income distribution. For example, if one sorts individuals into income quintiles on the basis of lifetime income using standard mortality rates, and then independently sorts them again using differential mortality rates, one finds that over 98 percent of individuals are in the same income quintile under either definition. The finding that differential mortality rates do not have a first order effect is also confirmed in the work of (Harris & Sabelhaus 2005). We have confirmed that this lack of a large mortality effect applies under other definitions of progressivity and income as well. Thus, in remaining rows, we will 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. Doing so means that many individuals who have low earnings but high wage rates (e.g., the college educated stay-at-home parent) will now be 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 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, one out of every eight (12.5%) of the individuals who were in the lowest quintile based on individual earnings are in one of the top three income quintiles defined based on potential income.

Because 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), the Gini coefficient (row 3) is lower on both a before and after tax basis. The before-tax Gini falls to 0.31, while the after tax Gini falls to 0.299. The EP measure declines to 1.016, suggesting that Social Security, while still slightly redistributive, is clearly less progressive when evaluated on the basis of potential earnings.

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

Row 4 of table 1 reports the result using 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 is substantially less redistribution than when progressivity is based on individual income measures, and it reflects the fact that much of the apparent redistribution from Social Security is happening within, rather than between, households. Indeed, the EP suggests that there is very little net redistribution from Social Security, as the presence of Social Security barely changes the Gini.

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. While these two cases do overlap – for example, they are both ways of increasing the measured well-being of high ability individuals who opt out of the labor force – these measures are not perfectly correlated. By combining both, we simultaneously recognize that households share resources *and* may optimally 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 and after Gini coefficients to 0.277 and 0.275, respectively. The EP falls to only 1.0029.

In row 6, we repeat this same case, but this time 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 virtually no effect on the overall level of income inequality when evaluated using a higher discount rate.

The overall conclusion from the analysis of Gini coefficients and effective progression is

that there is virtually no (or at best, very little) redistribution resulting from Social Security. What little redistribution there appears to be when evaluating the system based on individual, lifetime earnings nearly disappears when one considers 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.

6.2 But Does Social Security Help the Poor? Lifetime Tax Rates by Income Quintile

As noted earlier, the EP measure is designed to characterize the degree of progressivity across the entire income distribution. However, an equally legitimate policy concern is the extent to which Social Security does or does not help those 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?

In table 2, we explore the same cases as in table 1, only this time we report, by income quintile, the median net lifetime tax rate. 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 rate, indicating that these households have received 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.

Looking first at Table 2, row 1, we see that in the lowest income quintile, the median lifetime net tax rate from Social Security is *minus* 21%. This indicates that in this quintile, the median effect of Social Security was to *increase* after-tax lifetime earnings by 21%. In the next

lowest quintile, the median net tax rate was minus 10%. 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 as we move up 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.

Moving to the definition of potential income (row 3), the median tax rate in the bottom quintile is now -2.7% . This reflects the fact that, for a given low income individual, the denominator is now larger for having replaced low actual earnings with higher potential earnings. It also reflects the fact that the new definition has changed the composition of who is in the bottom quintile. Again, as one moves up to higher income quintiles, the median tax rate is increasing. 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 the potential income measure.

Using household, rather than individual, 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% . It is interesting, however, to note that there is also a change in the relative treatment of the upper two income groups when evaluated on a household basis. Specifically, the median tax rate in the 4th decile is actually somewhat *higher* than the median tax rate in the top decile – in other words, the system is regressive at the upper end of the income distribution. This could easily happen 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). This can also occur because of the regressive nature of the earnings cap, which limits the exposure of high earners to the payroll tax. We also see from Table 3 a further decline in the

precision with which net transfers are targeted to those in the bottom of the income distribution.

When one uses potential household income (row 5), one finds that the median lifetime net tax rate is now just slightly positive at 0.2%. 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, while 12-16% of those in the upper three income quintiles receive such transfers. Thus, while Social Security may, on average, transfer some resources to the lowest income quintile, the program is not well targeted in that it both a) fails to redistribute to a large fraction of those most in need, and b) unnecessarily transfers resources to large numbers of higher income individuals.

As expected, the use of a higher discount rate in row 6 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). 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.

A comparison of the net tax rate results with the effective progression results allows one to develop a better understanding of the underlying dynamics. The key feature is that, even when the measure of effective progression indicates very little redistribution, it is still the case that the individuals in lowest income quintile have, on average, significantly lower lifetime net tax rate from Social Security than do individuals in higher income quintiles. In the upper half of the income distribution, however, there system exhibits little progressivity, 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 progressive when one uses broader income measures.

Together, these results suggest three main conclusions. First, once one accounts for within household resource sharing as well as voluntary time allocation to household production

and/or consumption of leisure, the progressivity of Social Security 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, there are still many low income households that pay net taxes and many high income households that 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 more efficiently target resources to the lifetime poor is an interesting question for future research.

6.3 Is the Degree of Progressivity Changing?

The previous section demonstrated that a shift from actual to potential earnings, or a shift from individual to household resources, strongly influences the degree of measured progressivity 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 should be expected to change along with labor force participation patterns. 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 there are fewer one-earner couples and more two-earner couples, the extent of redistribution in the system might plausibly be expected to change.

To address this, we have split our PSID sample into two sub-samples based on their birth dates. The first sample is our "pre-Boomer" cohort, namely, those who were born prior to 1946. The second is our "Baby Boomer" sample, namely those born in 1946 or after. By repeating the analysis of section 6 on these two sub-samples, we can learn to what extent the different labor market attachments of these two groups influence the degree of progressivity.

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 6 scenarios. Using the traditional concepts of individual lifetime income, it would appear that Social Security is becoming less progressive due to the fact that the EP measure is larger for the Pre-Boomer generation than for the Baby Boom generation. When 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. Within the baby boomer sample, the EP is positive, although extremely small. The analysis of lifetime tax rates by income quintile and the analysis of the fraction of negative lifetime tax rates by quintile tell a similar story.¹⁷ Namely, the perceived decline in overall system progressivity that appears on an individual level as one compares pre-boomers and boomers is reversed when one examines the data using broader income concepts.

Overall, however, the examination of progressivity within a generation tells a very similar story to that of the overall sample. Namely, when using the most inclusive concept of income that accounts for the earnings potential of oneself and one's spouse, the Social Security system does not appear to be progressive in any meaningful way.

7. Conclusion

In this paper we seek to measure the extent to which the current social security system redistributes resources from rich to 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 when progressivity is measured using comprehensive concepts of income, the Social Security system exhibits less overall progressivity than when it is evaluated using more narrow definitions of income. For example, when evaluated using potential labor

¹⁷ In the interest of space, these tables are not included here, but are available from the authors upon request.

earnings at the household level (rather than actual individual earnings), the Social Security retirement program exhibits virtually no overall progressivity (as measured by the change in the Gini coefficient). Second, we find that result is largely driven by the lack of progressivity (and occasionally, the presence of regressivity) in the middle and upper part of the income distribution, whereas those in the bottom income quintile may, in fact, still benefit from the program in a progressive way. Third, we find that even when there is redistribution occurring, 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 extent to which progressivity differs across cohorts depends on the income concept used.

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 progressivity. Second, this analysis is purely a financial one, and thus misses the important insurance aspects of the Social Security program. Analysis of the insurance value of Social Security, including its providing of earnings insurance, longevity insurance, and disability insurance, 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 Progressivity 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 Progressivity 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 Progressivity?" 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 Progressivity 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 Progressivity." *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.

Table 1
Effective Progression

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

Table 2
Median Lifetime Net Tax Rate by Income Quintile

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

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

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

Table 4**Effective Progression Pre-Boomers vs. Baby Boomers**

Scenario	Lifetime Earnings Measure	Lifetime Earnings Basis	Mortality	Discount Rate	Effective Progression Full Sample	Effective Progression Pre-Boomer Sample	Effective Progression Baby Boomer Sample
1	Actual	Individual	Standard	2%	1.0315	1.0370	1.0286
2	Actual	Individual	Differential	2%	1.0309	1.0373	1.0276
3	Potential	Individual	Differential	2%	1.0159	1.0181	1.0146
4	Actual	Household	Differential	2%	1.0036	0.9991	1.0070
5	Potential	Household	Differential	2%	1.0029	1.0003	1.0047
6	Potential	Household	Differential	4%	1.0006	0.9974	1.0025