

Retirement Wealth Across Cohorts: The Role of Earnings Inequality and Pension Changes

Ann Huff Stevens



Project #: UM08-18

Retirement Wealth Across Cohorts: The Role of Earnings Inequality and Pension Changes

Ann Huff Stevens
University of California, Davis

September 2008

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
<http://www.mrrc.isr.umich.edu/>
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5-04). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Regents of the University of Michigan

Julia Donovan Darrow, Ann Arbor; Laurence B. Deitch, Bingham Farms; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; S. Martin Taylor, Grosse Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Retirement Wealth Across Cohorts: The Role of Earnings Inequality and Pension Changes

Ann Huff Stevens

Abstract

Changes in labor markets over the past 30 years suggest upcoming changes in the distribution of wealth at retirement. Baby boom cohorts have spent the majority of their prime earnings years in a labor market with increased earnings inequality. This paper investigates how changes in lifetime earnings distributions affect the distribution of retirement wealth among cohorts retiring over the next decade. I use data from the Health and Retirement Study from 1992 to 2004 to estimate the relationship between lifetime earnings, pre-retirement private wealth and Social Security wealth. I show that changes in the lower half of the male earnings distribution explain a substantial portion of changes in the distribution of pre-retirement wealth. Growth in women's earnings across the cohorts do not offset these declines in wealth associated with male earnings. When pensions are added to the measure of wealth, the role of earnings is even larger, reflecting a strong correlation between changes in earnings across these cohorts and changes in the values of their employer-provided pensions. These pension changes do not appear to operate via changes in pension structures (defined benefit versus defined contribution). The present value of wealth from future Social Security benefits, in contrast, grows in real terms throughout most of the distribution. At the bottom of the male distribution of Social Security wealth, reductions in lifetime earnings limit this growth in real benefits, while at the top of the distribution earnings growth amplifies expected growth in Social Security wealth.

Authors' Acknowledgements

Thanks to Jason Lindo and Rebecca Reed-Arthurs for outstanding research assistance.

Wealth and income at retirement are the result of many forces that accumulate over the lifetime: labor force participation and outcomes, public policies, savings decisions, and chance. As a result, long-lasting changes in any of these forces may affect the retirement wealth of cohorts with a substantial lag. In this spirit, major changes in the United States' labor market over the past 30 years predict changes in the determinants of wealth and an increase in the dispersion of wealth for cohorts now on the verge of retirement. The distribution of wages has undergone a transformation, with income inequality growing dramatically starting in the late 1970s. Because the cohorts who faced this increased earnings dispersion throughout the bulk of their prime earnings years have not yet reached retirement age, rising earnings inequality has not yet been clearly linked to inequality in retirement wealth. This study investigates how changes in labor market returns during the 1970s, 80s, and 90s translate into differences in the level and distribution of retirement wealth for birth cohorts now approaching retirement. Specifically, it compares the distribution of wealth across two cohorts of individuals, those born between 1936 and 1941, and those born between 1948 and 1953. Additionally, it examines the extent to which these changes in the distribution of retirement wealth can be linked to changes in the structure of earnings over time.

I show that the changing distribution of earnings can explain a substantial portion of the increased dispersion in wealth and the decline at the bottom of the distribution of pre-retirement wealth. In particular, changes in current and lifetime earnings across the two cohorts account for up to half of observed changes in the fractions of individuals with low asset levels. Changes in the value of pension benefits are highly correlated with

changes in earnings across these cohorts so that those with declining earnings have even larger declines in wealth when pension values are included in the wealth measures.

When Social Security wealth is considered, benefit calculation formulas, especially the indexing of benefits to average wage levels, mean that the real dollar value of average benefit levels rise across cohorts. The exception to this is at the lowest part of the wealth distribution for men, where rising real benefits for a given earnings level and falling real earnings have roughly offsetting effects on expected wealth from Social Security.

A large body of literature relates to the determinants, adequacy and distribution of retirement wealth, including work by Venti and Wise (1998), and Dynan, Skinner, and Zeldes (2004), among many others. Venti and Wise (1998), show that the vast majority of the dispersion in retirement wealth is unexplained by measures of lifetime earnings. The current study differs in focus, however, by concentrating on changes over time in the distribution of retirement wealth. It is possible, for example, that changes in the earnings structure will still explain a substantial portion of changes in retirement wealth. Dynan, Skinner, and Zeldes (2004), in contrast, argue that much variation in savings is related to earnings. They argue that measurement error in earnings, particularly near the bottom of the distribution, could drive some of the apparent non-responsiveness of wealth to observed income levels. Dynan, Skinner, and Zeldes further argue that savings rates increase with permanent income. These conclusions are consistent with a role for changes in the wage distribution in explaining the evolution of retirement wealth. This role could be limited, however, by very low savings rates among those with the lowest wages, where much of the change in the wage distribution has occurred.

Much work has also centered on documenting and explaining changes over time in the distribution of wealth. James Smith (1999), for example, examines a number of potential explanations for rising wealth inequality between the early 1980s and late 1990s. He rejects the possible contribution of income inequality to wealth inequality and argues instead that the distribution of capital gains and the behavior of equity markets is the main explanation. Smith, however, does not directly test the empirical relationship between changing income inequality and changing wealth inequality. Scholz (2003) also investigates wealth inequality using an approach that relies on comparisons across cohorts and reaches a slightly different, though tentative conclusion concerning the role of income inequality. He compares the evolution of wealth inequality, particularly at the top of the distribution, to that of income inequality and concludes that “similar factors may be behind trends in both series.”

Another segment of the literature has looked specifically at retirement wealth among recent cohorts leading up to the baby-boomers who are the subject of this study. Karen Smith (2003) and Bosworth, Burtless, and Sahm (2001) address the question of how changing earnings inequality may affect retirement income across cohorts and how recent cohorts will differ from earlier generations of retirees. Neither of these studies, however, directly examines the non-pension wealth accumulation or details of private pension wealth across cohorts.

This study makes several contributions to the literature on retirement wealth. First, I use the most recent data available from the Health and Retirement Study (HRS) on the wealth of the baby-boom cohort as of 2004, when they are in their early 50s. Wealth as of age 50 is a strong predictor of wealth at retirement, and so this will provide

the most up-to-date forecast of how this cohort is likely to fare in retirement. Second, I directly address the relationship between changes in the distribution of earnings and changing wealth distributions by relating measures of lifetime earnings to measures of current wealth. While there has been a large amount of attention devoted to documenting and explaining the details and causes of evolving wage inequality, this study offers a somewhat unique look at one of the consequences of this wage inequality as it is transmitted into wealth holdings near the end of individuals' working lives. Third, my focus on the early baby-boom cohort (and the earlier cohort born twelve years prior), is particularly relevant for an investigation of the role of earnings inequality. Because the major rise in earnings inequality began at the end of the 1970s, and because more experienced workers benefited initially from rising returns to experience, the full impact of rising earnings inequality on retirement wealth may only recently be fully evident. The later cohort I study was between ages 30 and 50 during the years from 1978 through 2003, so that virtually all of their prime working years occurred under a regime of expanded inequality in earnings. While the earlier cohort used here (who were 30 to 50 years old between 1966 and 1991) were not entirely unaffected by rising earnings inequality, they are substantially less exposed to the major changes in the wage structure. This is the earliest cohort surveyed by the HRS.¹ Finally, this paper will estimate the association between changing earnings distributions and changing distributions of wealth in a way that can identify differential changes at different points in the wealth distribution. I use the decomposition methods developed and applied by Fortin and

¹ It would be useful to have an even earlier cohort for comparison purposes. I attempted to make use of data from the Retirement History Survey cohort interviewed in 1969, along with information on their earnings histories. Differences in the wealth data, however, made it impossible to make direct comparisons across these surveys.

Lemiux (1998) to relate changes in the distribution of wages across cohorts to changes at different parts of the distribution of wealth.

An important caveat to this approach is that is based on the observed association, across cohorts, between lifetime earnings and wealth. The usual difficulty in reaching conclusions about causality, involving unobservables that may be correlated with both earnings levels and savings or investment behavior, will apply to the results here. Thus, these results can establish a link between rising earnings inequality and rising wealth inequality, but cannot clearly establish that all of this link runs from earnings themselves (rather than associated, but unobserved, factors) to wealth accumulation.

I. Data and empirical approach

A. HRS Data

Data from this study come from Waves 1 through 7 of the Health and Retirement Study (HRS). The sample consists of two groups of respondents who were born between 1936 and 1941 (cohort 1) and between 1948 and 1953 (cohort 2). Individuals in cohort 1 are initially observed in wave 1 of the HRS, collected in 1992 when they are 51 to 56 years old, with follow-up information on wealth closer to retirement taken from waves 6 and 7. Individuals in cohort 2 are first observed in wave 7 of the HRS, as a part of the newly introduced “early boomers” cohort. To be consistent with cohort 1 observations, individuals observed in wave 7 of the HRS are included if they were ages 51 to 56 at the time of the 2004 survey. I chose this age range primarily based on data availability, but these cohorts may have ten or more years to accumulate wealth before retirement. I refer to the measures as pre-retirement wealth, and in section II.B. below I include results that

use later measures of wealth for cohort 1 and predict later wealth distributions for cohort 2.

For most of the analysis, I also require data on total household wealth, subsets of household wealth (such as non-housing wealth), and individual lifetime earnings histories. Many elements of household wealth in the HRS.. are imputed to some degree. For this analysis, I drop all individuals who do not report at least partial bracketed information for all asset measures. I focus on two primary measures of household wealth, total wealth, and non-housing wealth. Total wealth includes all real estate, business, and financial assets (stocks, mutual funds, checking and savings accounts), IRA or Keogh accounts, and all debts including mortgages on the primary residence. Non-housing wealth drops the net value of the primary residence from this measure. Wealth resulting from entitlement to future Social Security benefits is calculated, making use of the administrative data on lifetime earnings histories to calculate expected future Social Security benefits for both cohorts of individuals.

In some of the analysis, I also incorporate self-reported data on expected wealth from employer-provided pensions. The self-reported data are not ideal for predicting actual retirement wealth, since it has been well-documented (see, for example, Gustman and Steinmeier, 2004 and Chan & Stevens 2008) that many individuals do not accurately report their pension wealth. Unfortunately, while employer-reported pension data are available for the earlier cohort I study, such data are not yet fully available for the later cohort.

To predict the connection between wealth and an individual's earnings, I would ideally include the entire lifetime earnings history. Here, I make use of summary

earnings records from the Social Security Administration, linked to HRS respondents, and available to researchers under certain restrictions. These records contain total earnings from W-2 records for each year from 1978 through 2003, for those respondents who, in 2004, gave the HRS permission to obtain their wage records. The primary advantage of these data is that they provide a rich earnings history over many years for the respondents for whom they are available. An important disadvantage of these earnings data is that not all HRS respondents gave permission for these records to be obtained. Nearly 50 percent of my initial cohorts do not match to the lifetime earnings histories. This is a lower match rate than has been reported in earlier work using an earlier set of Social Security earnings histories collected at wave 1. Earlier work has shown few significant differences between HRS respondents who do and do not have a match to the restricted earnings history files. (Haider and Solon, 2000). Using the 2004 permissions file, I similarly find few differences in observables between those who do and do not provide permission for this link.

Appendix Figures 1a and 1b compare the distribution of household wealth and current income among all respondents in these two cohorts, and those who match to the restricted earnings files. The distributions correspond reasonably well, but there are statistically significant differences in several of the variable means. Mean asset holdings are significantly lower in the restricted sample, while educational attainment is statistically significantly higher in the restricted sample. Given these differences between the full and matched samples, some caution is necessary in applying the findings below, all of which are based on the restricted sample.

Lifetime earnings used in the analysis take these earnings records and average them over ages 36 to 51 for each cohort. This allows for consistent treatment of both cohorts, while making use of as many years of earnings data as possible.

B. Empirical Approach

The goal of the empirical work here is to link increasing dispersion in lifetime earnings across these two cohorts to potential increases in wealth dispersion as the cohorts approach retirement. To capture changes at different parts of the wealth distribution, and reflect changes in that distribution itself requires an alternative to simple regression-based decomposition methods. My approach follows that used by Fortin and Lemieux (1998) who examine changes in the male-female wage gap at various points throughout the wage distribution.² Specifically, to capture changes in wealth across the distribution, I divide the distribution of wealth into a number of intervals, $a_j, j = 1, \dots, J$. I then estimate an ordered probit to predict probabilities of an individual having wealth that falls into any one of these individual intervals. Thus, for each cohort ($c=1,2$), I can calculate a set of predicted probabilities of having wealth in J different intervals. These probabilities are calculated as $\hat{\pi}_j^c$ where

$$(1) \hat{\pi}_j^c = \frac{1}{N_c} \sum \{ \Phi(\hat{\lambda}_j^c - X_i \hat{\beta}^c) - \Phi(\hat{\lambda}_{j-1}^c - X_i \hat{\beta}^c) \}$$

and β and λ are the coefficients and thresholds from the ordered probit estimation, and the summation is over all the individuals in cohort c .

² An alternative, though largely equivalent, approach is that by Machado and Mata (2005). This approach requires fewer parametric assumptions than the ordered probit method used here. Given the relatively small size of my data for each cohort, the probit method seems preferable, and I have used it throughout.

Once the empirical probabilities are calculated, it is straightforward to calculate counterfactual probabilities of the distribution of individuals across wealth categories that would prevail with different Xs, (or different β s). In particular, below I calculate the counterfactual probabilities that individuals in cohort 2 would have had wealth in the given interval, if they had faced the same distribution of lifetime earnings (and other Xs) as cohort 1. This is calculated as

$$(2) \hat{\pi}_j^{2,cf} = \frac{1}{N_1} \sum \{ \Phi(\hat{\lambda}_j^2 - X_i \hat{\beta}^2) - \Phi(\hat{\lambda}_{j-1}^2 - X_i \hat{\beta}^2) \}$$

where the summation now occurs over all Xs in cohort 1.

II. Results

A. Summary statistics

I begin by presenting summaries of the distribution of wealth, earnings and other characteristics across the two cohorts. In these summary statistics, and throughout the analysis, I focus on measures of wealth per person, in which I divide household wealth by two for individuals currently married or cohabitating. While I control for marital status in many of the results below, using wealth measures that are adjusted in this way for marital status reduces the importance of changes across the cohorts in marital status and histories, to better focus on earnings distribution changes.³ I also estimate the relationship between earnings and wealth separately for men and women throughout most of the analysis, despite most of the asset measures being collected at the household level.

³ For analysis of the importance of changes in demographics, particularly marital status see recent work by Mammen (2008) and Zissimopoulos, Rauer, and Karney (2008).

I do this because male and female earnings distributions have evolved quite differently over time, and separating men and women in this way makes it easier to isolate these differential changes in earnings distributions.

Table 1 begins with summary statistics across the two cohorts for measures of wealth, earnings and demographic indicators. Statistics reported below are based on tabulations using the household sample weights from the HRS. Beginning with the distribution of non-pension, private wealth, Table 1 shows a small increase in mean wealth for men, decreases at the mean for women, and reductions in median wealth for both genders. When the expected present value of pension wealth is included in the measure of wealth, the across-cohort changes in wealth are similar. Mean wealth now increases, and median wealth declines for both men and women. Mean and median Social Security wealth increases in real terms for both men and women across cohorts, resulting primarily from the use of wage indexing in the Social Security benefit formula which keeps benefits rising over time in real terms.

Lifetime earnings for both cohorts are measured as average annual earnings between ages 36 and 51. Cohort 1's lifetime earnings are averaged over the calendar years 1978 through 1992, and cohort 2's are averaged over 1989 through 2003. Lifetime earnings also move in the expected ways. For men, there is a 13 percent decline in median real earnings, reflecting the deterioration of earnings power for relatively low-skilled men over this period. For women, in contrast, there are substantial increases in average earnings at both the mean and median.

Other important changes across these two cohorts are summarized in Table 1 as well. The structure of pension benefits has changed over time, with the frequency of

defined benefit pensions declining substantially among men, and women. When both defined benefit (DB) and defined contribution (DC) pensions are considered there is a reduction in the probability of having a pension for men and a slight increase for women. Individuals in cohort 2 have higher mean household income, are more educated, and are less likely to be currently married.

Because my main interest here is in the dispersion in wealth and earnings, Table 2 summarizes the distributions of wealth and lifetime earnings. The main facts to note from Table 2 are the large dispersion in household wealth for both cohorts in these data and the increase in this dispersion among cohort 2. Among men, the ratio of wealth at the 25th percentile of the distribution to that of the 75th is approximately .17 among the earlier cohort, and falls to .07 in the later cohort. Among women, dispersion also increases, with the ratio of wealth at the 25th to 75th percentile falling from .14 to .07. The increase in women's wealth in cohort 2 occurs only at points above the 75th percentile. For both men and women, wealth declines for the more recent cohort at all points of the distribution at or below the median. Median private wealth falls by approximately 25 percent across these cohorts separated by just 12 years. Other authors report generally consistent trends in wealth across these cohorts. Lusardi and Mitchell (2007), for example, show declines in total net worth below the median, and increases above the median when they compare the same cohorts from the HRS. Using the Survey of Consumer Finances from 1983 to 2001, Wolff (2007) reports increasing wealth inequality over time among those aged 45 and over.

The next panel of Table 2 shows the distribution of wealth excluding the value of the primary residence, and highlights the extent to which much savings occurs through

individual's investment in their homes. At the median of the total wealth distribution, eliminating equity in the primary residence reduces the asset measure by more than fifty percent. While the level of wealth when housing equity is excluded is much lower, patterns across the distribution and across cohorts are similar to the patterns for total wealth. One difference is that up to the 75th percentile for both men and women, the level of non-housing wealth has fallen for cohort 2. Another difference from total wealth is that non-housing wealth shows no increase for women even at the 90th percentile.

While pension data for cohort 2 in this survey are somewhat imperfect, and come only from self-reports of pension values, I have calculated total wealth for these cohorts, including the value of (self-reported) employer-provided pensions. For individuals reporting defined contribution pensions, I simply take the reported account balances, in 1992 dollars, as of the survey date for each individual in the two cohorts. For individuals reporting defined benefit pensions, I take the reported benefit amount expected if the individuals retires as the "normal" retirement age, and calculate the present value of these benefits from that age to age 100, adjusted for the probability of survival using actuarial tables from the Social Security Administration.

Pension plus non-pension wealth distributions are shown in the third section of Table 2. The addition of pension wealth substantially increases the level of wealth, with the added mass of the distribution concentrated above the median of non-pension wealth. At the median, men's wealth including pension values increases to \$97,000, and women's to \$70,000. Below the median, there are only minor increases in wealth, since relatively few of those with low non-pension wealth have substantial wealth from pension plans. It is important to note that my use of the self-reported data, particularly in this

simple tabulation, is likely to understate total pension wealth. For the distribution shown in Table 2, individuals who report having a pension but do not report information on the value of that pension are assigned zero pension wealth. In regression results below I control for these missing values.

To begin exploring the link between dispersion in retirement wealth and dispersion in lifetime earnings, the final panel of Table 2 shows the distribution of average lifetime earnings for the two cohorts. Patterns here confirm the well-known increase in dispersion in earnings for men over this time period. Echoing the patterns in the lower half of the wealth distributions shown above, lifetime earnings among men fall substantially across cohorts at points up to and including the median. Male earnings are roughly stable at the 75th percentile, and increase significantly over time at the 90th percentile. The distribution of earnings among women evolved quite differently over this time period, with earnings (and participation) rising throughout the distribution. At the 25th percentile, for example, women's earnings rose from approximately \$6000 for the first cohort to approximately \$8000 for the second cohort. At the 75th percentile, earnings also rose substantially, from \$21,000 to \$24,000 per year.

B. Linking Earnings and Private Non-pension Wealth

While the general pattern of changes in earnings and wealth across these cohorts suggests a connection between the two, it is important to quantify this relationship, given the relatively weak connection often noted between earnings and wealth, particularly among low earners. I next estimate ordered probit models, following the technique used in Fortin and Lemiux (1998). The idea is to use a number of intervals of the wealth

distribution so that observables can influence probabilities of being located at various points in the distribution of wealth. I begin by dividing the distribution of wealth into 20 discrete intervals. The choice of 20 intervals is arbitrary, but reflects two goals. First, the advantage of this approach is that it allows flexibility in the relationship between wealth and lifetime earnings (and other regressors) over the full distribution of wealth.⁴ If too few intervals are chosen, the full, and potentially non-linear, relationship will not be accurately captured. Second, because I have a relatively small data set, I must limit the intervals to have an adequate number of observations in each interval. The intervals are set based on dividing the cohort 1 male distribution of wealth into 20 equal quantiles. These quantiles are calculated using unweighted data, to ensure a sufficient (and equal) number of observations in each interval.

To predict the distribution of wealth from lifetime earnings, I begin with a specification that includes only the earnings variables as regressors. Separately for each cohort, I estimate the probability of wealth in each interval as a function of 4th order polynomials in average lifetime earnings, and in current year earnings. I continue to separate men and women, and use wealth measures reflecting household totals normalized by the number of adults in the household. This allows for a separate focus on how the different changes in the male and female earnings distributions are related to the evolution of these wealth distributions. In later specifications I add spousal earnings, and a handful of other observables, to the specifications.

The results of these initial regressions are generally as expected. The coefficient estimates are not immediately meaningful in the ordered probit setup, but one key feature of these results should be highlighted. The estimated coefficients relating lifetime

⁴ I have also estimated these models allowing 40 intervals of wealth and obtained very similar results.

earnings to wealth are very similar across the two cohorts. This suggests that there is little role for changes in the coefficients linking earnings and wealth in explaining changes in retirement wealth. Changes in the coefficient on lifetime earnings might suggest changes across these cohorts in savings propensities, but I find no evidence consistent with this possibility. The lack of differences between the two cohorts in these key coefficients motivates my focus on changes in the earnings variables (as opposed to changes in coefficients) in explaining changes across the cohorts. Most of the change in wealth accumulation not explained by the earnings and other variables can thus be attributed to unobserved factors.

Figure 1 shows histograms that summarize the distributions of total wealth among the twenty specified intervals for the two cohorts. The histograms show both the empirical frequencies of wealth in each interval and the corresponding predicted probabilities based on the ordered probit estimation. Echoing the summary statistics above, the figures illustrate the increased dispersion in wealth among cohort 2. The fraction of wealth concentrated at both the ends of the distribution is substantially larger for cohort 2 than for cohort 1, particularly at the low end of the distribution. The figure also shows that the predicted probabilities capture the actual probabilities and their movement across cohorts quite well.

Figures 2a and 2b are based on the same estimation, but present the cumulative fractions of men and women with wealth below the level indicated on the X-axes. The lines without markers in the figures show the cumulative distribution of pre-retirement wealth for cohorts 1 and 2. For men, the story of increased dispersion is illustrated by the crossing of the lines representing the two cohorts at per-person wealth holdings of

approximately \$50,000, near the median of the distribution. The left-hand portion of the figure shows that individuals from cohort 1 are less likely to have relatively low asset accumulations. As we move to the right in Figure 2a, however, the line for cohort 1 is above that for cohort 2, showing that those in cohort 1 are also less likely to have very *high* wealth. Among women, the cohort 1 line is again below that for cohort 2 at points throughout most of the distribution consistent with the distributions summarized in Table 2. I discuss the “counterfactual” lines in these figures below.

Next, I ask the extent to which individual’s lifetime earnings histories can explain these changes in the distribution of wealth across cohorts. Specifically, as own earnings inequality has increased, has inequality in pre-retirement wealth followed suit? Table 3 begins to answer to this question. To summarize changes across the distribution of wealth, I pick several, fixed (in 1992 dollars) points in the distribution of non-pension wealth, and summarize the fraction of individuals above or below that point in each of the two cohorts. For consistency across the cohorts and across different segments of Table 3 I use the cohort 1 male (unweighted) distribution to define the cutoff points throughout the table. Specifically, the table shows the fraction with wealth below the 25th percentile, below the median, and above the 75th percentile.⁵

The first line of the table shows that 23 percent of men from cohort 1 had pre-retirement wealth of less than \$19,300. Among cohort 2, this fraction grew to almost one-third, or 32 percent. Forty-seven percent of men in cohort 1 had wealth below \$53,000, compared to 54 percent in the later cohort. At the top of the distribution, 27

⁵ Note that, because the intervals were chosen to equalize the number of observations in each interval, they are based on the unweighted distribution, but results are based on weighted regressions and tabulations. Thus, there are, for example, 23% of wave 1 men with assets below the (unweighted) 25th percentile for that same group.

percent of cohort 1 men and 29 percent of cohort 2 have wealth holding in excess of \$115,000. Among women, the comparable figures are 29 percent for cohort 1 and 26 percent for cohort 2. Thus, the early baby boomers are substantially more likely to have very low stocks of pre-retirement wealth than the earlier cohort. At the top of the distribution, there is a small increase in the fraction of men with assets more than \$115,000.

Next, I calculate the counterfactual distribution of wealth for cohort 2 if they had faced the same distribution of lifetime and current earnings as cohort 1. These calculations are summarized in columns 3 and 6 of Table 3. Changes in male earnings explain a substantial fraction of the change in the wealth distribution across these cohorts. The counterfactual prediction for the fraction of cohort 2 men with wealth below \$19,300 falls to 29 percent and the fraction with less than \$53,000 falls to 51 percent. When these men are assigned lifetime earnings of the earlier birth cohort, out of a 9 percentage point change in the fraction with very low wealth, 3 percentage points, or 33 percent of the total change, is accounted for. The fraction of men with wealth below \$53,000 increases by 7 percentage points, approximately half of which is explained by changes in the earnings distribution.

Returning to Figure 2, the line with round markers shows the counterfactual distribution of wealth throughout the distribution. The fraction of the across-cohort change in the distribution of wealth that is explained by earnings changes increases through the bottom half of the wealth distribution. At levels slightly above median wealth, the change in earnings across these cohorts can account for all of the change in wealth. Given the relationship between earnings and wealth captured by the ordered

probits, if men in cohort 2 had faced the same distribution of earnings as men in cohort 1 there would be no difference in the fractions of those cohorts with wealth below approximately \$62,000 at ages 51 to 56.

Among women, the role of earnings changes is starkly different and quite limited. Since women's earnings increased throughout the earnings distribution between the two cohorts, we would expect earnings changes to predict increased wealth, which did not occur. Indeed, the counterfactual exercise suggests that, if cohort 2 women had the (lower) earnings of cohort 1, there would have been a slightly higher fraction (one percentage point) of women with low household wealth, relative to what actually happened for cohort 2. The implied differences between the cohort 2 distribution and the counterfactuals are quite small throughout Table 3 for women. It is also important to note the remarkably low explanatory power that women's lifetime earnings have in the probits for wealth accumulation, particularly among cohort 1. The pseudo-R squared measure for the ordered probits for cohort 1 women, where the only independent variables are polynomials in women's own lifetime average earnings and current earnings is only .01. For comparison, the R-squared from the comparable probit for men is .04. In both cases even fairly rich measures of lifetime earnings explain relatively little of the variation in wealth, but for women this connection is particularly weak.

Of course, for most women in these cohorts, spousal earnings have made up a larger fraction of household income over their lifetimes than their own earnings. For both men and women, any story linking household wealth accumulation to changing earnings distributions should focus not only on the distribution of respondents' own earnings, but also on the distribution of earnings for spouses. Rising average earnings among women

over the lifetimes of these two birth cohorts should offset some of the deterioration of male earnings at the bottom of the distribution since most of the men are (or have been) married. This is reflected in Table 1 (sample means) by the increase in real household income across cohorts from \$64,000 to \$68,000 for men and \$52,000 to \$59,000 for women. In panel 3a of Table 3 I show the results of repeating the earlier multinomial probit estimation, but including additional explanatory variables to capture total earnings in the household. Specifically, I add to the probit specification of wealth intervals an indicator for whether there are two, married or cohabitating adults in the household, and interactions between the married indicator and spouse's average lifetime earnings and current earnings. In this case, the counterfactual simulation assigns the cohort 1 values of both own and spousal lifetime and current earnings, along with marital status, to cohort 2.

The initially surprising result is that, when spouse's earnings are included as one of the Xs in the probit estimation for males, there is virtually no change in the extent to which these observable characteristics explain the increase in the fraction of cohort 2 men with wealth. This is surprising since rising women's earnings should offset some of the reduction in male earnings for these birth cohorts. The explanation for this, however, lies in the necessity of controlling for current marital status when I include spousal earnings. Panel 2 of Table 3 reports an intermediate step, in which I add controls (beyond own earnings) only for current marital status. The addition of this covariate raises the fractions of changes in wealth below \$19,300 and \$53,000 that are accounted for by observables (which now include earnings and marital status) to between 40 and 90 percent. Because fewer men in cohort 2 are currently married, and because being married is associated with higher wealth holdings (even on a per-person basis), a simple indicator

for marital status explains a substantial fraction of the change. In panel 3, where I also include spousal wages for those who are married, the fraction explained by observables again falls, reflecting that the rising earnings of spouses of the married men in these cohorts did offset some of the effects of their own earnings changes on wealth.

Figure 3 summarizes the results of the counterfactual exercise for men's total non-pension wealth, using both own and spousal earnings to predict wealth. The black bars show the total change in the fraction of men with wealth below the amount indicated on the horizontal axis. The white bars show what this change would be if there had been no change in the distribution of earnings across cohorts. Up to the 55th percentile, holding constant earnings across cohorts reduces the increase in the fraction with low wealth by 20 to 50 % of the total change. Above the median, earnings explain none of the shift in the wealth distribution, since holding earnings constant would have resulted in even larger increases in the fraction with wealth below the indicated amounts. A final point to note is that the changes in the estimated wealth distribution above the median are substantially smaller than those at the bottom of the distribution, and are often not statistically different from zero.

When I focus on results for women, but include their spouse's lifetime and current earnings as one of the observable predictors of pre-retirement wealth, there continues to be little contribution of changing Xs to the overall change in their wealth distribution. Including husband's earnings changes the counterfactual analysis only slightly.

As mentioned in the discussion of Table 2 above, much of the accumulation of wealth for these age groups involves housing equity. This may be important to take account of because, if much of the value of these assets is related to appreciation in the

housing stock, rather than savings or investment behavior, it would not be surprising if earnings were unrelated to this portion of wealth. Specifically, we might expect earnings to be a more powerful indicator of non-housing wealth, since this measure of wealth cannot be directly affected by changes in housing values. As panel 3b of Table 3 shows, however, the results are similar when housing equity is eliminated from the measure of wealth. In particular, among women, own earnings continue to explain virtually none of the change in wealth holdings. Because the average level of wealth is substantially reduced when housing equity is subtracted, I also report the sizeable fractions of individuals with non-housing wealth below \$8,700 in this section of the table. This also captures a more comparable segment of the non-housing wealth distribution. Nearly one-third of cohort 1 men fall into this category, and that fraction increases to 39 percent among cohort 2. Changes across the cohorts in lifetime and current earnings can account for almost 40 percent of the change in the fraction of men with virtually no assets beyond their primary residence as of ages 51 to 56.

Before moving to other components of wealth, it is worth considering what these measures of wealth at ages 51 to 56 will translate into at ages closer to actual retirement. I focus on the age 51 to 56 range because that is the oldest age at which I can currently observe cohort 2. Using data for cohort 1, however, it is possible to see what the distribution of wealth looks like 10 years later, at wave 6, when they are aged 61 to 66. Table 4 begins by presenting simple tabulations of the distribution of wealth for cohort 1 taken from wave 6 of the HRS. As expected, wealth has grown substantially over the decade since this cohort was first observed, with median wealth increasing by 35 percent among men and 38 percent among women relative to their wave 1 distribution.

Using the relationship between wave 1 and wave 6 wealth estimated for the first cohort allows for the calculation of predicted wealth levels for cohort 2 under several different assumptions. First, I estimate the relationship between the level of wealth held by cohort 1 as of wave 6 as a function of the level of wealth held at wave 1. To link these results to the analysis thus far, I use the same ordered probit approach to predict the probability of wealth from the wave 6 data in each of 20 intervals. Now, however, the key independent variable is wealth in wave 1. Table 5 summarizes some points in the resulting distribution of predicted wealth. Next, I use the coefficients from these probits, combined with the initial wealth of cohort 2 (from the 2004 data) to predict what their distribution of wealth will look like ten years later. This prediction, of course, assumes that cohort 2's wealth will evolve in the same way that cohort 1's wealth evolved between 1992 and 2002. The first two columns of Table 5 summarize this exercise, which not surprisingly, looks very much like the results in Table 3. This simply inflates the wealth of cohort 2 to the level that might be expected when they are closer to retirement. Twenty-nine percent of this cohort is predicted to reach their 60s with non-pension wealth of less than \$26,800.

To capture the spirit of the counterfactual exercises above, I can also predict cohort 2's wealth at ages 61 to 66 under different assumptions about their lifetime earnings. This is more speculative than the counterfactual exercises above, since I am holding constant both the relationship between earnings and wealth at ages 51 to 56, and the evolution of this wealth over the next ten years. As a first step I predict wealth at ages 51 to 56 for cohort 2 based on their lifetime earnings information. In step two, I use the cohort 1 coefficients (relating wealth at ages 61 to 66 to the earlier wealth measures)

to predict wealth for cohort 2 at ages 61 to 66. Column 3 of Table 5 shows these results. The first thing to note is that the distribution of wealth in column 3 is less disperse, reflecting that these predictions do not use any previous wealth measures as a predictor. I do this to have a basis for comparison when I recalculate this predicted wealth distribution holding earnings constant at the distribution of cohort 1. Because most of the variation in wealth at ages 51 to 56 is unrelated to earnings, there is less spread in these predicted values based only on earnings. For this reason, direct comparisons should be made only between columns 3 and 4 in this table. Finally, in column 4, I repeat the exercise just described, but assign the cohort 1 earnings levels to cohort 2 and predict their wealth. This isolates the effect of changing earnings when compared with results in column 3. Again, changing earnings predict substantial changes in the distribution of wealth. Column 4 results for men suggest that just 20 percent of these men would be predicted to have wealth below \$26,800 at retirement ages if they had not faced the changing earnings distribution that prevailed over their working years. This is compared to a prediction of 26 percent with wealth below this level given the earnings they actually experienced.

C. Linking Earnings to Total (Pension and Non-pension) wealth

The discussion of private wealth for retirement thus far is incomplete because it ignores the potentially important role of employer-provided pensions plans. Because many individuals can anticipate income flows during retirement from defined benefit pensions, or have assets accumulated in defined contribution pension plans, non-pension wealth may seriously understate their overall resources available during retirement. As

noted above, I do not currently have employer-reported data on pensions available for cohort 2, and so I must rely on self-reported pension details to construct measures of pension wealth for both cohorts.

Table 6 repeats the analysis of Table 3, but adds the present value of pension wealth to the measure of per-person wealth from the earlier analysis. Note that the dollar values for the intervals summarized in Table 6 differ from the earlier table, but again capture roughly the same points in the distribution, which has shifted to the right with the addition of pension wealth. The fraction of men with this broader measure of wealth of less than \$24,000 goes from .23 among cohort 1 to .30 among cohort 2. The counterfactual assigning cohort 1 earnings to cohort 2 reverses almost all of this change. Looking at the results for wealth below \$82,500, the counterfactual suggests that earnings changes more than account for the decline in wealth between cohorts. These results should be interpreted carefully. Reduced earnings could clearly translate into reductions in non-pension wealth, holding constant savings and investment behavior over the lifetime. The link between falling earnings below the median and reduced employer-based pension wealth must operate through a different mechanism. The story here simply suggests highly correlated changes in earnings and pensions for individuals in these cohorts. In particular, this pattern suggests that those men for whom real wages were falling most dramatically have also seen reductions in the value of their pensions. This is consistent with work by Pierce (2001) who shows that compensation inequality increased somewhat more than wage inequality between 1981 and 1997. Pierce also shows that pensions are an important contributor to this rise in compensation inequality.

Figure 4 illustrates the role of earnings changes throughout the distribution of total private wealth of men. At points in the distribution below approximately \$70,000, most or all of the increase in fractions with wealth less than a given cutoff would be reversed if earnings had not declined. Implicitly, this also requires that the relationship between earnings and non-wage compensation (in particular, pensions) is also held constant.

One particularly discrete change in pension plans over time has been the switch from employers offering defined benefit pensions to defined contribution plans. Table 7 shows the fraction of each cohort of men and women who report having a defined benefit pension from a current employer. As expected, the fraction with a defined benefit plan rises as lifetime earnings rise. The percentage decrease in frequency of DB plans across cohorts is also somewhat correlated with individuals' place in the wage distribution. Men in the bottom quartile of earnings have seen a reduction of more than 50 percent in their likelihood of having a defined benefit pension, while those in the top quartile have seen a reduction of around 38 percent.

To explore whether this movement away from defined benefit plans can account for any of the change in wealth distributions, beyond that already accounted for by earnings changes, I add an indicator for having a defined benefit pension plan to the probit estimation. As seen in the lower part of Table 6, this changes the results based on own and spousal earnings only slightly. What is clear from Table 6 is that much of the decline in pre-retirement wealth holdings below the median is associated with the decline in earnings below the median across the two cohorts studied here. These results suggest that this occurs both because individuals with lower earnings have accumulated less savings,

and because these lower earnings were accompanied by lower levels of employer-provided pension wealth.

D. Linking Earnings to Social Security Wealth

While private wealth accumulation is of interest both as a reflection of resources for retirement and because it reflects a behavioral response to altered earnings environments, wealth from Social Security benefits provides, for many households, the bulk of potential retirement savings. In this section, I turn to examination of how changes in earnings across these cohorts translate into changes in the present discounted value of lifetime Social Security benefits. For each respondent in the restricted sample used in the tables above, I also have information on covered Social Security earnings from 1951 to 2003, and quarters of work from 1947 through 2003. There are two differences between these records and those used to estimate the lifetime earnings measures used above. First, these records are top-coded at the Social Security earnings taxable maximum, so that these reflect earnings actually used to estimate benefits. Second, these earnings cover a longer portion of the respondents' working lives, again to allow more accurate forecasting of benefits. These records are used to calculate the respondents' expected Social Security benefits, assuming initial claiming at the normal retirement age. I utilize information on both own and spousal earnings from these records to calculate expected benefits. I also assume that the current spouse survives until age 65. In calculating the present value of benefits, I assign survival probabilities for each year from the current survey year (1992 or 2004) that vary with both age and gender, taken from the Social Security Administration's life tables. I assume a 3% real rate of discount, and then calculate the

present value of benefits to be received, where the discounting is done as of 1992 or 2004, depending on the cohort.

The potential role for earnings changes here is very different than in the analysis of private household wealth. Here the relationship is driven primarily by a mechanical, non-linear transformation of average lifetime earnings (up to the Social Security taxable maximum in each year) into benefit amounts. More specifically, for each individual, the top 35 years of earnings are chosen, and earlier years are indexed to the earnings level in the year the individual turns 60. Because these cohorts are observed several years prior to their eligibility for benefits, I assume that they continue to earn the same annual amount (in real terms) and that their work hours remain relatively constant. To qualify for Social Security benefits, individuals must work for at least 10 years. I assume that individuals who are currently working as of age 51 to 56, when they are observed, will continue to accrue quarters of service at approximately the same rate, and so assume that their years of service at age 65 will be 1.3 times that observed at approximately age 53. This is likely to slightly overstate eligibility for benefits, but the vast majority of individuals in these samples qualify for benefits even with less generous assumptions.

For those who have worked for at least 10 years, calculation of Social Security benefits first involves the calculation of “average indexed monthly earnings”, or AIME. This involves taking the highest 35 years of earnings, indexing them to the level of wages as of the year the individual turns 60, and averaging the indexed earnings. Next, the benefit formula is related to the AIME using a three part schedule, with two specified “bend-points” that establish the fraction of earnings to be replaced by benefits. For

example, cohort 1 will reach age 65, on average, in the year 2004. In that year, the basic benefit amount (known as the “primary insurance amount”) is calculated as:

$$\begin{aligned} & (.90 * AIME) \text{ up to } \$612 \text{ in AIME} + \\ & (.32 * (AIME - \$612)) \text{ up to } \$3689 \text{ in AIME} + \\ & (.15 * (AIME \text{ above } 3689)). \end{aligned}$$

Bend points (\$612, and \$3689 in the formula above) are indexed to the general level of wage increases in the economy across years.

I assume, as a simplification, that the current spouse is the only current or former spouse under whom each individual may have some benefit eligibility. This ignores the possibility that some individuals will qualify for different benefit amounts as the result of a previous marriage lasting 10 years or more.

Table 8 summarizes the distribution of Social Security wealth across the two cohorts, and shows substantial growth in benefit levels between the two cohorts. The primary reason for this growth over time is related to the benefit formulas and the method of indexing. As is documented in Biggs, Brown, and Springstead (2005, hereafter BBS), the current method of using wage indexing to hold replacement rates roughly constant at the mean, results in substantial growth in the (price-indexed) real level of benefits across cohorts. They show that, with no changes to the Social Security benefit calculation formulas, the average level of benefits rises by more than 10 percent each decade. Thus, the increase across cohorts reflects the fact that Social Security formulas are indexed to average wage growth, and so increase at a faster rate than the general price level. As we will see below, this feature of Social Security guarantees that benefits (and thus Social

Security wealth) rise at the median, where real wages, and thus AIME, are relatively constant across cohorts for men, and rise substantially for women.

Table 8 also summarizes the distribution of Average Indexed Monthly Earnings (AIME), a key component in the calculation of benefits. Movements in AIME should, of course, be highly correlated with movements in lifetime earnings reported in Table 2, but AIME is included here because it is the direct input into the calculation of Social Security benefits. The 25th percentile of the distribution of AIME among men declines by approximately 15 percent, from \$2290 to \$1908 across cohorts. The 25th percentile of per capita social security wealth, in contrast, grows slightly. For women, changes in the distribution of AIME reflect growth in the earnings levels of women across these cohorts. Women's total social security wealth grows more slowly than earnings over this period, reflecting the fact that, for most women, their benefit eligibility associated with their husband's earnings continues to dominate benefit levels based on their own earnings.⁶ Median Social Security wealth for women is projected to grow by roughly 25 percent from cohort 1 to cohort 2.

Changes in the distribution of Social Security wealth based on the probit decomposition that links total social security wealth to lifetime earnings are summarized in Table 9. Among men, the fraction with Social Security wealth of less than \$147,000 is constant at 0.23 in both cohorts. The ongoing increase in real Social Security benefits (for a given level of lifetime earnings, or AIME) exactly offsets the decline in lifetime real earnings at this point in the distribution. The counterfactual exercise summarized in

⁶ Married individuals are entitled to either their own calculated Social Security benefits, or 50% of the level of benefits based on their spouse's earnings, whichever is greater.

the third column confirms this. If cohort 2 had faced the same earnings distribution as cohort 1, the fraction with Social Security wealth below \$147,000 would fall to 0.18.

Moving further up the distribution of Social Security wealth, approximately half of the men in cohort 1 are projected to have Social Security wealth below \$188,000. Among cohort 2 this fraction falls to 0.39. The final column for men shows that, holding earnings constant, this increase in benefits would have been more pronounced, with fewer than one-third expecting Social Security wealth below \$188,000.

On the right hand side of Table 9, the analysis is repeated for women. Among women, the combination of wage-indexing of benefits and rising real earnings produces large increases in expected benefit levels throughout the distribution of Social Security wealth. The counterfactual exercise decomposes these two factors and shows that, holding earnings levels constant, the proportions of women with Social Security wealth below \$147,000 (\$188,000) would have fallen to .38 (.55). The bottom panel of Table 7 shows that, adding marital status and spousal earnings as observables to be held constant in computing the counterfactual distribution has relatively little impact.

All of the calculations of Social Security wealth presented thus far assume that the benefit structures and formulas remain unchanged over the course of cohort 2 (and cohort 1's) remaining working and retirement lives. This is unlikely given the future fiscal difficulties facing the Social Security program. In particular, the current schedule of real benefits increasing faster than the rate of price increases seems unlikely to be maintained indefinitely. As BBS (2005), point out, several alternatives to the wage-indexing approach used to calculate current benefits could move the Social Security program closer to fiscal sustainability. To illustrate this effect, and to demonstrate what Social

Security wealth might look like if aspects of the benefit formula are changed, Table 10 summarizes results that recalculate benefits under one of the potential scenarios studied in BBS (2005). These authors note that, if the bend points in the Social Security formula were adjusted according to a price index, rather than an index of wage growth, the growth across cohorts in average benefits would be sharply curtailed. The 1st two columns of Table 10 simply repeat results from Table 9 for comparison, while the 3rd column shows projected Social Security wealth assuming that the bend points in the Social Security formula were indexed according to the CPI starting in 2004. For men, this produces a distribution of retirement wealth that is closer to stable across the two cohorts. In particular at the bottom of the distribution, the fraction of men with a present value of benefits from Social Security below \$147,000 would increase only slightly, from .23 to .25 under the altered benefit calculation formula. At higher benefit levels, there would continue to be growth in benefits across cohorts, but this growth would be reduced. Among women, benefit levels would continue to grow in real terms throughout the distribution of benefits.

Two key conclusions emerge from the analysis of Social Security wealth in this context. First, the expected increases in average benefits levels that will result if Social Security benefit formulas remain unchanged serve to counteract the effects of changes in the earnings distribution on Social Security wealth. At the lowest part of the earnings distribution, the growth generated by the wage indexing of Social Security benefits just counteracts the decline in real wages experienced by men in these cohorts. While Social Security is widely acknowledged to play an important role in reducing poverty among the elderly, this highlights another distributional contribution of the Social Security program;

Social Security will play an important role in counteracting the decline in lifetime earnings for relatively unskilled workers and maintaining a fairly constant level of retirement wealth at the bottom of the wealth distribution. At points higher in the distribution there is an increase in Social Security wealth from both constant or rising wages, and normal growth in benefits. Second, the analysis comparing wage- and price-indexing shows that this particular modification to benefit calculation rules results in fairly small increases in the fraction with low Social Security wealth. Other reforms, of course, could produce different effects across the distribution. Given the substantial changes at the bottom of the private wealth distribution among individuals who will soon retire, the role of Social Security for individuals in this part of the wealth distribution may be an important consideration in evaluating potential program changes.

III. Conclusions

These findings suggest that the changing distribution of earnings may play an important role in explaining changes in the level of private retirement wealth across recent and future cohorts of retirees. At some level it is not surprising to find that those individuals whose real earnings have declined substantially will arrive at retirement with fewer assets accumulated. On the other hand, low rates of savings among low income earners might have been expected to minimize the impact of these reductions on wealth. Work by Venti and Wise (1998) has suggested that much of the dispersion in retirement wealth is the result of choices concerning how much to save out of income over the lifecycle. These results suggest the same is not true when we focus on how wealth levels are changing across these cohorts. Reduced earnings over several decades, particularly at

the bottom of the wealth distribution, have produced sizeable changes in accumulated wealth.

In contrast to the findings for private, pre-retirement wealth, Social Security benefits are not predicted to decline over time, even for those groups who have faced significant deterioration in their real earnings. The results here highlight the importance of considering how potential reforms to Social Security benefit formulas might affect the entire distribution of benefits. For those groups whose lifetime earnings were lower than previous cohorts, the role of Social Security in their overall wealth portfolio may be larger than for earlier cohorts.

References

- Biggs, Andrew, Jeffrey R. Brown, and Glenn Springstead. 2005. "Alternative Methods of Price Indexing Social Security: Implications for Benefits and System Financing" NBER Working paper #11406.
- Bosworth, Barry, Gary Burtless, and Claudia Sahm. 2001. "The Trend in Lifetime Earnings Inequality and Its Impact on the Distribution of Retirement Income." Washington: The Brookings Institution.
- Chan, Sewin, and Ann Huff Stevens. 2008. "What You Don't Know Can't Help You: Pension Knowledge and Retirement Decision-Making." , *Review of Economics and Statistics*.
- Dynan, Karen, Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy*, 112:397-444.
- Engen, Eric, William Gale, and Cori Uccello. 1999. "The Adequacy of Household Saving." *Brookings Papers on Economic Activity*, 2:65-187.
- Fortin, Nicole and Thomas Lemieux. 1998. "Rank Regressions, Wage Distributions, and the Gender Gap." *Journal of Human Resources*, 33(4):610-643.
- Gustman, Alan L., and Thomas L. Steinmeier. 2004. "What People Don't Know about their Pensions and Social Security" in William G. Gale, John B. Shoven and Mark J. Warshawsky, editors, *Private Pensions and Public Policies*, (Washington D.C.: Brookings Institution).
- Gottschalk, Peter, and Robert Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers on Economic Activity*, 2:217-272
- Haider, Steven J. 2001. "Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991." *Journal of Labor Economics*, 19:799-836.
- Lusardi, Annamaria, and Olivia S. Mitchell. 2007. "Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth." *Journal of Monetary Economics*, January 54(1): 205-224.
- Machado, Jose, and Jose Mata. 2005. "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression." *Journal of Applied Econometrics* 20:445-465.

- Mammen, Kristin. 2008. "The Long-Term Effects of the Divorce Revolution: Health, Wealth, and Labor Supply." Working paper prepared for 2008 Annual Joint Conference of the Retirement Research Consortium.
- Scholz, John Karl. 2003. "Wealth Inequality and the Wealth of Cohorts." Mimeo, University of Wisconsin.
- Smith, James. 2001. "Why is Wealth Inequality Rising?" In *The Causes and Consequences of Increasing Inequality*. Ed. Finis Welch.
- Smith, Karen. 2003. "How Will Recent Patterns of Earnings Inequality Affect Future Retirement Incomes?" Washington: The Urban Institute.
- Venti, Stephen, and David Wise. 1998. "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *American Economic Review*, 88:195-191.
- Wolff, Edward N. 2007. "The Retirement Wealth of the Baby Boom Generation." *Journal of Monetary Economics*, January 54(1):1-40.
- Zissimopoulos, Julie, Benjamin Karney, and Amy J. Rauer. 2008. "Marital Histories and Economic Well-Being." Working paper prepared for 2008 Annual Joint Conference of the Retirement Research Consortium.

Figure 1 Actual and Predicted Fractions of Cohort with of Wealth in Each Interval

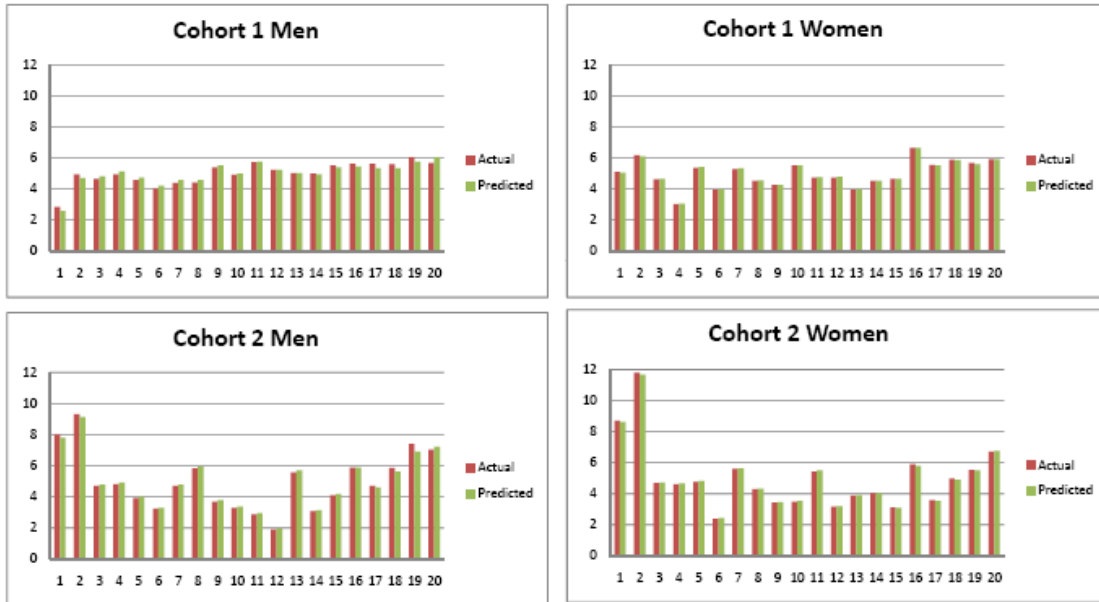


Figure 2a
Men: Fraction with Wealth Below Indicated Level

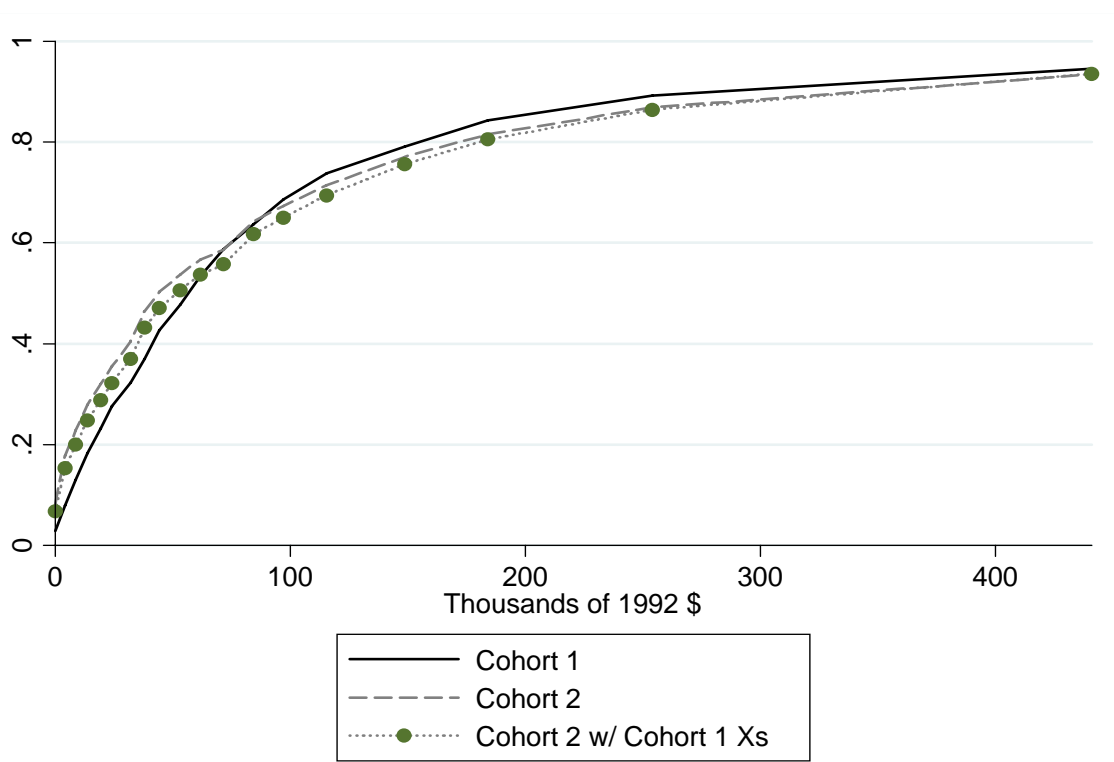


Figure 2b
Women: Fraction with Wealth Below Indicated Level

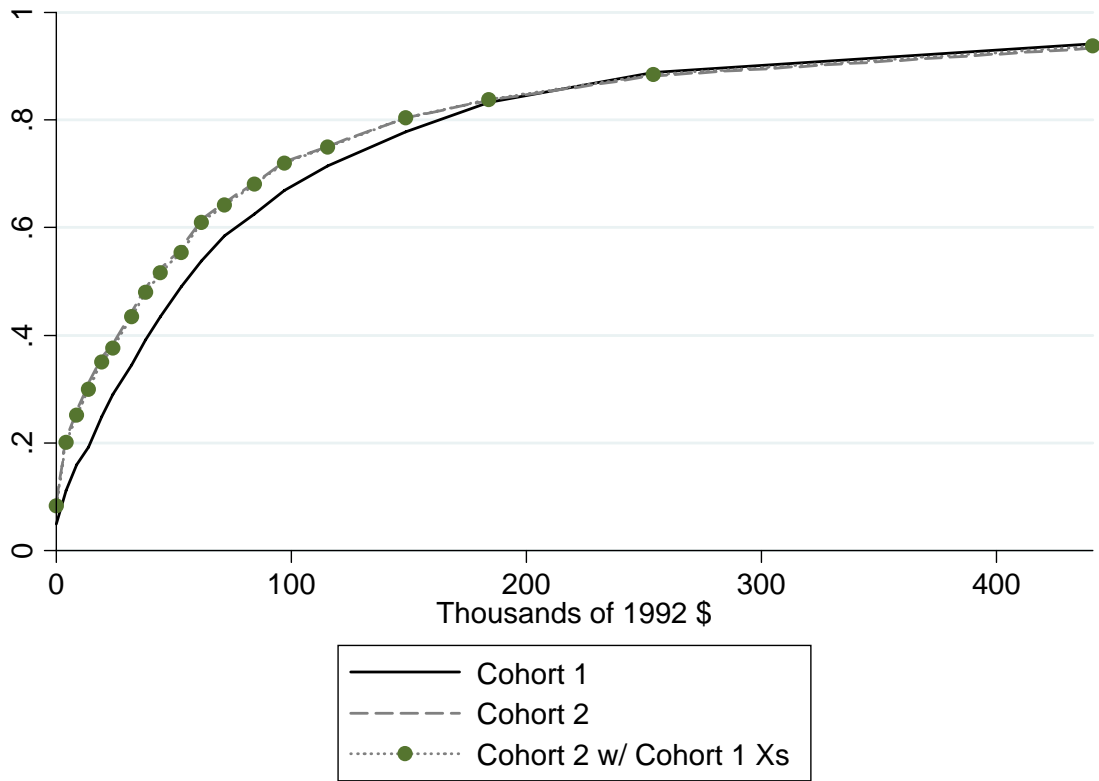


Figure 3

Change in Cumulative Distribution of Wealth: Total & Earnings Constant (Men)

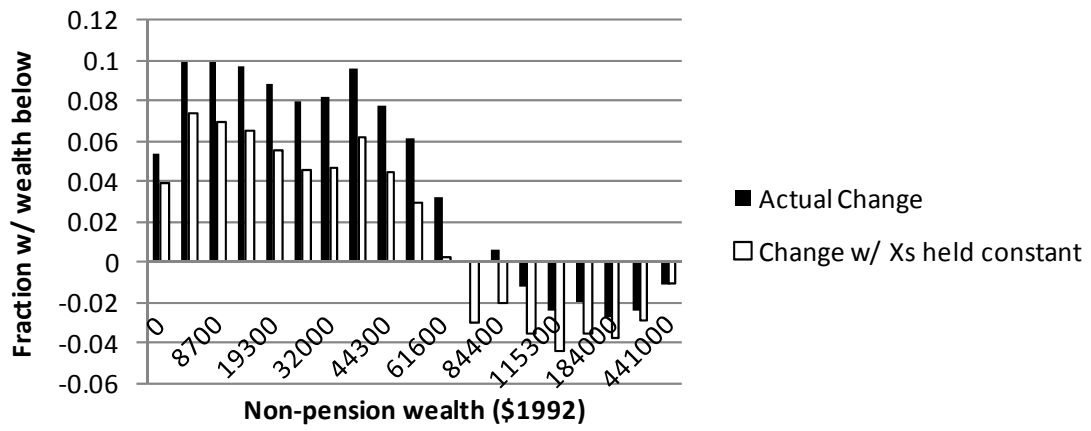


Figure 4

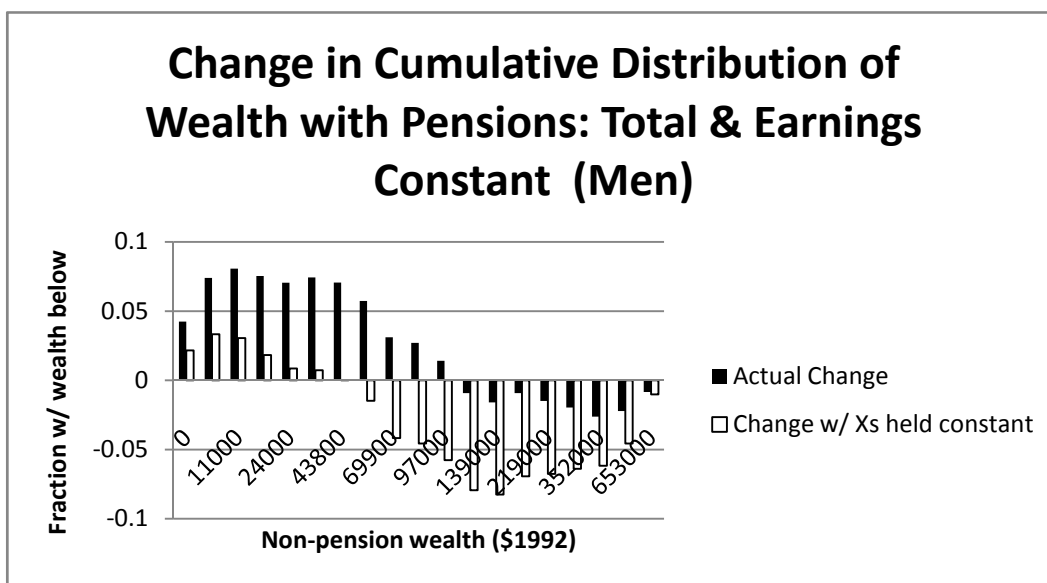


Table 1
Sample Characteristics

	Men		Women	
	Cohort 1	Cohort 2	Cohort 1	Cohort 2
Non-pension private wealth per person (\$1992)				
Mean	130,437	134,707	125,741	111,715
Median	60,450	47,734	57,000	43,796
Non-pension plus pension private wealth per person (\$1992)				
Mean	200,753	225,688	148,882	160,846
Median	96,864	87,110	69,849	60,549
Social Security Wealth per person (\$1992)				
Mean	188,588	211,760	152,962	191,843
Median	192,342	221,269	156,461	196,059
Lifetime Average Earnings (\$1992)				
Mean	28,316	27,605	14,715	17,531
Median	29,232	25,295	12,724	14,889
Defined Benefit Pension	0.37	0.20	0.27	0.20
Any Pension	0.58	0.50	0.44	0.45
Total Household Income	63,992	68,799	51,695	58,531
Years of Education	13.1	13.6	12.8	13.6
Married	0.85	0.72	0.71	0.59
# Times Divorced	0.48	0.62	0.46	0.79
N=	1040	617	1038	554

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004. Data are from waves 1 and 7 of the HRS, weighted using household level weights.

Table 2
Distribution of Wealth and Lifetime Earnings

	Men		Women	
	Cohort 1	Cohort 2	Cohort 1	Cohort 2
Non-pension private wealth per person (\$1992)				
Percentile				
10th	5,000	0	2,623	0
25th	23,000	11,114	19,850	8,543
50th	60,450	47,734	57,000	43,796
75th	135,150	149,703	137,000	129,272
90th	275,000	355,126	282,500	342,867
Non-pension, non-housing, private wealth per person (\$1992)				
Percentile				
10th	1,800	-371	258	-1,857
25th	6,750	2,417	5,000	965
50th	26,500	18,276	24,500	12,630
75th	80,500	78,009	80,500	56,463
90th	212,500	258,172	212,500	167,719
Non-pension plus pension private wealth per person (\$1992)				
Percentile				
10th	7,400	334	3,809	118
25th	33,500	18,172	24,050	12,258
50th	96,864	87,109	69,849	60,549
75th	245,208	266,063	168,500	187,480
90th	475,372	512,630	338,670	413,819
Lifetime Average Annual Earnings				
Percentile				
10th	2,358	4,287	2,539	3,557
25th	15,403	11,869	6,243	8,285
50th	29,232	25,295	12,724	14,889
75th	42,880	42,487	20,973	24,391
90th	49,552	56,042	30,177	35,665

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004. Data are from waves 1 and 7 of the HRS, weighted using household level weights.

Table 3
Contribution of Observables to Changing Distributions of Non-pension Private Wealth

		Men			Women		
		Cohort 1	Cohort 2	Cohort 2 with: Cohort 1 Xs	Cohort 1	Cohort 2	Cohort 2 with: Cohort 1 Xs
Included Xs: Own Earnings							
1	Total Assets	(1)	(2)	(3)	(4)	(5)	(6)
	Wealth < \$19,300	23%	32%	29%	25%	35%	36%
	Wealth < \$53,000	47%	54%	51%	48%	54%	56%
	Wealth > \$115,000	27%	29%	31%	29%	26%	25%
Included Xs: Own Earnings & Marital Status							
2	Total Assets						
	Wealth < \$19,300	23%	32%	28%	25%	36%	35%
	Wealth < \$53,000	48%	54%	49%	49%	55%	55%
	Wealth > \$115,000	27%	29%	31%	29%	26%	25%
Included Xs: Own Earnings, Marital Status & Spouse Earnings							
3a	Total Assets						
	Wealth < \$19,300	23%	32%	29%	25%	36%	35%
	Wealth < \$53,000	48%	54%	51%	49%	56%	55%
	Wealth > \$115,000	26%	29%	31%	29%	25%	25%
3b	Non-housing assets						
	Wealth < \$8700	31%	39%	36%	33%	45%	43%
	Wealth < \$19,300	46%	53%	50%	47%	59%	58%
	Wealth < \$53,000	68%	72%	71%	70%	76%	75%
	Wealth > \$115,000	17%	17%	17%	18%	14%	14%

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004.
Predicted probabilities of wealth in given range based on ordered probit estimation using household level weights.

Table 4
Distribution of Wealth at Ages 61-66 for Cohort 1
(\$1992)

	Men	Women
	Cohort 1	Cohort 1
Household Non-pension Wealth		
Percentile		
10th	7,018	3,898
25th	33,612	29,245
50th	85,787	78,768
75th	223,825	204,718
90th	455,839	417,235

Note: Cohort 1 are born 1936 to 1941 and observed in 1992. Data are from wave 6 of the HRS, weighted using household level weights.

Table 5
Alternative Predictions of Cohort 2 Wealth at Ages 61-66

	Men			
	1	2	3	4
	Cohort 1	Cohort 2	Cohort 2	Cohort 2
Wealth predicted by:	Wealth at ages 51-56	Wealth at ages 51-56	Cohort 2 lifetime earnings	Cohort 1 lifetime earnings
Predicted non-pension wealth				
Wealth < \$26,800	23%	29%	26%	20%
Wealth < \$73,000	50%	54%	55%	51%
Wealth > \$195,000	25%	24%	19%	18%
	Women			
	Cohort 1	Cohort 2	Cohort 2	Cohort 2
Wealth predicted by:	Wealth at ages 51-56	Wealth at ages 51-56	Cohort 2 lifetime earnings	Cohort 1 lifetime earnings
Predicted non-pension wealth				
Wealth < \$26,800	25%	29%	22%	19%
Wealth < \$73,000	52%	55%	51%	47%
Wealth > \$195,000	22%	22%	19%	20%

Note: Cohort 1 is born 1936 to 1941 and observed in 1992 and 2002; cohort 2 is born 1948 to 1953 and observed in 2004.

Table 6
Contribution of Observables to Changing Distributions of Pension plus Non-pension Private Wealth

	Men			Women		
	Cohort 1	Cohort 2	Cohort 2 with:	Cohort 1	Cohort 2	Cohort 2 with:
			Cohort 1 Xs			Cohort 1 Xs
Included Xs: Own Earnings, Marital Status & Spouse Earnings						
Wealth < \$24,000	23%	30%	24%	26%	34%	34%
Wealth < \$82,500	48%	51%	45%	55%	60%	61%
Wealth > \$219,000	25%	27%	30%	17%	20%	19%
Included Xs: Own Earnings, Marital Status, Spouse Earnings & DB Pension Flag						
Wealth < \$24,000	23%	30%	23%	25%	33%	33%
Wealth < \$82,500	48%	51%	44%	56%	60%	59%
Wealth > \$219,000	25%	26%	32%	17%	20%	20%

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004.
 Predicted probabilities of wealth in given range based on ordered probit estimation using household level weights.

Table 7
Fraction with DB pensions by Lifetime Earnings Quartile

Lifetime Earnings Quartile	Men		Women	
	Cohort 1	Cohort 2	Cohort 1	Cohort 2
1	22%	10%	12%	17%
2	25%	13%	17%	12%
3	47%	31%	29%	17%
4	54%	33%	49%	31%

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004. Data are from waves 1 and 7 of the HRS, weighted using household level weights.

Table 8
Distribution of Projected Social Security Wealth and Average Indexed
Monthly Earnings (\$1992)

	Men		Women	
	Cohort 1	Cohort 2	Cohort 1	Cohort 2
Lifetime Social Security Wealth per Person				
Percentile				
10th	113,006	111,591	67,611	92,279
25th	154,818	161,517	107,267	131,971
50th	192,342	221,269	156,461	196,059
75th	222,654	262,085	196,317	243,852
90th	254,747	300,860	234,035	291,453
Average Indexed Monthly Earnings				
Percentile				
10th	1,315	978	301	504
25th	2,290	1,908	668	971
50th	3,349	3,356	1,283	1,865
75th	4,291	4,872	2,191	2,959
90th	5,184	6,252	3,181	4,264

Table 9
Contribution of Observables to Changing Distributions of Social Security Wealth

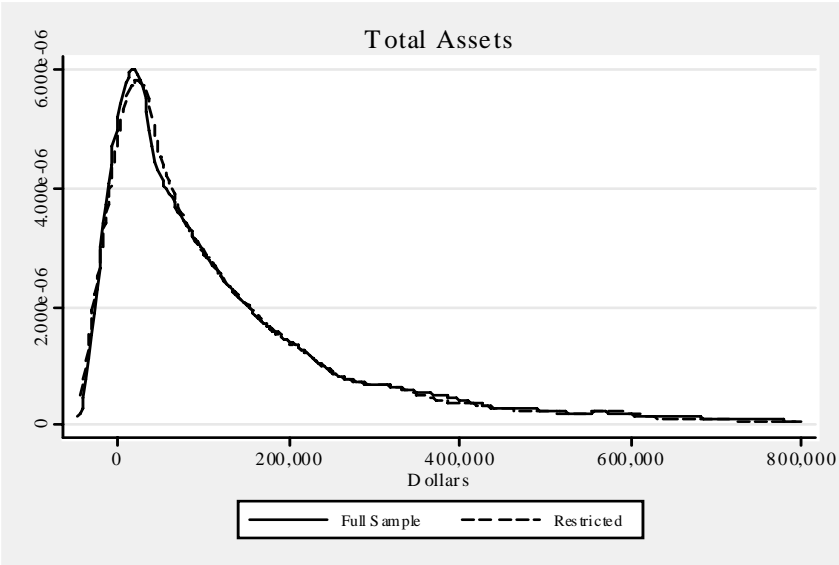
	Men			Women		
	Cohort 1	Cohort 2	<u>Cohort 2 with:</u> Cohort 1 Xs	Cohort 1	Cohort 2	<u>Cohort 2 with:</u> Cohort 1 Xs
Included Xs: Own Earnings						
SS Wealth per person < \$147,000	0.23	0.23	0.18	0.46	0.33	0.38
SS Wealth per person < \$188,000	0.49	0.39	0.32	0.70	0.50	0.55
Included Xs: Own Earnings, Marital Status & Spouse Earnings						
SS Wealth per person < \$147,000	0.23	0.23	0.21	0.47	0.32	0.37
SS Wealth per person < \$188,000	0.49	0.38	0.35	0.73	0.50	0.54

Table 10
Effects of Wage vs. Price Indexing Social Security Across Cohorts

	Wage Indexing		Price Indexing of
	Cohort 1	Cohort 2	Bend Points Cohort 2
Men			
SS Wealth per person < \$150,000	0.23	0.23	0.25
SS Wealth per person < \$200,000	0.49	0.38	0.45
Women			
SS Wealth per person < \$150,000	0.47	0.32	0.36
SS Wealth per person < \$200,000	0.73	0.50	0.57

Note: Cohort 1 is born 1936 to 1941 and observed in 1992; cohort 2 is born 1948 to 1953 and observed in 2004. Price indexing results assume that bend-points in the benefit formula are adjusted by the CPI.

Appendix Figure 1a
Distribution of Assets in Full and Matched Samples



Appendix Figure 1b
Distribution of Earnings in Full and Matched Samples

