University of
Michigan
Retirement
Research
Center

Working Paper
WP 2004-074

Modeling Lifetime Earnings Paths: Hypothetical versus Actual Workers

Andrew Au, Olivia S. Mitchell, and John W. R. Phillips

M R
R C

Project #: UM02-C2

"Modeling Lifetime Earnings Paths:
Hypothetical versus Actual Workers"

Andrew Au
University of Pennsylvania

Olivia S. Mitchell
University of Pennsylvania

John W. R. Phillips
Social Security Administration

March 2004

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104

Modeling Lifetime Earnings Paths:
Hypothetical versus Actual Workers

Andrew Au
Olivia S. Mitchell
John W. R. Phillips

Abstract

To assess the distributional effects of social security reform proposals, it is essential to have good information on real-world workers' lifetime earnings trajectories. Until recently, however, policymakers have relied on hypothetical earnings profiles for policy analysis. We use actual lifetime earnings data from the Health and Retirement Study (HRS) to compare actual workers' covered earnings profiles to these hypothetical profiles. We show that the hypothetical profiles do not track earnings patterns of current retirees; thus lifetime pay levels are much higher than for most HRS workers. Therefore, using hypothetical profiles could misrepresent benefits paid and taxes collected under such reforms.

Authors' Acknowledgements

Modeling Lifetime Earnings Paths:
Hypothetical versus Actual Workers

Understanding the distributional impact of social security benefits and taxes over the lifetime is a subject of growing interest in the policy arena. For instance, for the 1994-96 Social Security Advisory Council (ACSS 1996), the Office of the Chief Actuary (OACT) at Social Security assessed the likely distributional impacts of alternative program reforms according to their effects on three types of hypothetical earnings profiles. In 2001, the Commission to Strengthen Social Security requested analysis from the agency to better understand the impact of alternative reform scenarios on low, medium, and high earners (CSSS, 2001). And recently, the 2003 Technical Panel Report to the Social Security Advisory Board proposed that the Social Security Administration (SSA) generate an annual Benefit Audit offering "a comprehensive assessment of taxes paid and benefits received by different types of households and different groups of people." That report further expressed concern that the Trustees do not currently provide "an assessment of the net value of social security coverage and how this varies by household and personal characteristics" (Technical Panel, 2003: 15).

Without question, policymakers in the future will increasingly be asked to analyze the distributional implications of social security programs and possible reforms. In order to produce accurate distributional assessments of social security policy proposals, of course, it is critical to have good information on real-world workers' lifetime work and pay trajectories. Until recently, however, researchers have been limited in their ability to evaluate how benefit or tax changes might change actual workers' economic opportunities and constraints. This was because the system was established in a different technological era when data on workers' annual earnings over their entire lifetimes were not collected. Instead, partial information was initially gathered

on taxable earnings up to the annual social security earnings cap or taxable maximum, and in some cases only for specific years. Hence, for long periods of time, averages of actual earners' total labor earnings were not available for analysis. More recently, actual earnings reports have become available to the SSA (in collaboration with the Internal Revenue Service), so better data on total earnings are now available for the American workforce contributing to the social security system. Indeed, some researchers have looked to use these data on earnings to create a more diverse set of hypothetical workers for policy analysis (Bosworth et. al 1999). SSA and the Congressional Budget Office as well as academic analysts have made use of administrative earnings records matched to survey data to conduct distributional analysis of social security and retirement well being[1].

This paper explores a set of hypothetical earnings profiles frequently used for social security policy analysis, and it then compares these to "real-world" lifetime earnings trajectories of actual workers followed over time by the Health and Retirement Study (HRS). The real-world information on workers' lifetime earnings is used to compare how well *actual* workers' covered earnings profiles correspond to the *hypothetical* streams associated with "low", "medium", and "high" profiles in social security policy simulations. Our results suggest that the hypothetical profiles do not closely mimic those of actual workers' earnings: both the shape of the hypothetical earnings profiles and the lifetime average values are substantially higher than for most actual workers. Therefore, examining program changes geared toward improving solvency with hypothetical profiles could misrepresent both payroll tax collections and beneficiaries' well-being.

---

[1] See O'Hara, et. al (2004) for a summary of the long-term policy simulation model developed by the Congressional Budget Office. For an example of distributional work using the SSA Modeling Income in the Near Term (MINT) model, see Butrica et. al (2003). Gustman and Steinmeier (2001) use the HRS matched to administrative earnings to consider the degree of redistribution in the current law Social Security program.

Our discussion begins with a description of the national average earnings series used by SSA to convert workers' pay earned early in life into earnings levels comparable to pay received later in life. We then examine the demographic, social, and economic characteristics of lifetime earnings. Finally, we evaluate the implications of the work for retirement income policy. This research contributes to an understanding of how to model the distributional impacts of both current program rules and future policy changes under the social security system.

Social Security Earnings Patterns Over Time

The SSA summarizes earnings profiles in different ways, depending on the purpose of the exercise. Initially, the agency's main objective was to compute an Average Wage Index (AWI), traditionally used in the determination of retired worker benefits. Subsequently, hypothetical earnings profiles based on the AWI has been developed to address distributional questions: for example, Nichols et al. (2001) use these profiles to compare internal rates of return for low, medium, and high earning workers in the cross-section and over time, under the current social security program and under reform alternatives.[2] In what follows, we discuss how this index was created, and we also explore how this index is used to create "hypothetical profiles" of lifetime earnings. Finally, we detail the procedure used to construct earnings for scaled workers.

The Average Wage Index:[3] A retiring worker's social security benefit is computed as a fraction of that worker's lifetime *indexed* earnings, such that pay in all years is converted into

---

[2] Indeed, the profiles are employed for several types of distributional analyses as described by Nichols et. al (2001: np): "Hypothetical earnings histories have traditionally been used by the Office of the Chief Actuary to illustrate a range of benefit levels, replacement rates, money's worth measures, time to recover contributions, and internal rates of return under the Social Security program. These illustrations have long been used to evaluate the program under present law, but have been increasingly been used to evaluate the effects of possible alternatives."
[3] This section draws heavily on Myers (1993).

comparable terms for some final year.[4] This relies on an indexation process where each year's

earnings are adjusted by the AWI, so as to determine lifetime average pay used in computing

retirement benefits.

One way to convert earnings from one year into a different year's dollar value would be

to use a consumer price deflator to do the translation. This would convert nominal earnings

received in year t into the same *real* (inflation-corrected) pay level in some later year $t+i$,

maintaining the same real purchasing power of the dollars over time. This approach, however,

was not taken by the SSA. Instead, the agency "wage indexes" benefits. What this means is that

SSA computes the worker's Average Indexed Earnings by converting his pay earned prior to age

60 into "age 60" earnings, by multiplying his nominal earnings in all years t (<60) by the ratio of

the AWI in the year the worker turns age 60, to the AWI in year t. This type of adjustment

therefore converts his *relative* pay each year into the same *relative* pay level just prior to

retirement; the technique maintains a worker's *relative* pay standing over time as compared to

the national average. Since earnings tend to rise faster than inflation over time, the AWI-

indexing approach has the effect of boosting a worker's actual lifetime average earnings over

time *in real terms,* increasing it by the rate of growth in the AWI to age 60.

The AIME concept is therefore intentionally designed to rise in real terms over time, due

to national productivity increases as reflected in the AWI. For this reason, initial retiree benefit

payments tied to this index also grow over time: that is, the purchasing power of initial benefits

for a retiree who turned 60 in 1990 falls below that of an identical retiree who earned the same

---

[4] The 1977 Social Security Act proposed this method of indexing earnings records based on average national
earnings for 1977 and thereafter (cf Myers, 1993: 176).

real pay over his lifetime but who turned 60 ten years later.[5] To determine how important this divergence is, Figure 1 depicts changes in the AWI and the Consumer Price Index (CPI-U), where the latter is the usual inflation measure that tracks changes in a typical "basket" of consumer goods.[6] The first panel (A) shows changes in the indices themselves, while panel (B) demonstrates how the earnings adjustments associated with these two indices differ for a worker attaining age 60 in 2001. Over time, the AWI adjustment has been much larger than the inflation adjustment, and the gap is depicted by the dashed black line. Clearly, AWI indexing of earnings produces a much higher indexed earnings tally than would price-indexed earnings: thus, pay from the mid-1950s would be boosted almost 10-fold using the AWI approach, versus by less than 7 using the CPI.

Figure 1 here

Constructing the Average Wage Index:[7] In the early days of the social security program, it was difficult for analysts to acquire the earnings data necessary to develop the AWI. As a result, different approaches were taken at different times, as is detailed in several actuarial notes which we summarize in Figure 2.

Figure 2 here

1973-1977: The base period for the AWI is defined as the period 1973 to 1977. It is interesting that, some 30 years later, earnings data collected during this short and long-ago period remain the basis for the AWI used today. When it was first implemented, the AWI was generated using

---

[5] To bring the Social Security system to solvency, the Commission to Strengthen Social Security recommended that earnings be indexed according to prices rather than the AWI. This would imply that benefits are maintained in real terms but not boosted in real terms over time as has occurred for the last 25 years (CSSS, 2001).
[6] The Bureau of Labor Statistics tracks the prices of goods and services consumed by typical urban households. Prices for the goods and services used to calculate the CPI are collected in 87 urban areas throughout the country and from about 23,000 retail and service establishments. Data on rents are collected from about 50,000 landlords or tenants. The weight of each consumption item is based on household consumption data from the Consumer Expenditures Survey. See http://www.bls.gov/cpi/home.htm for more information on the CPI-U.
[7] This section draws heavily on Nichols et al. (2001).

only earnings collected by the agency during the first quarter of the year from the Summary Earnings Record (SER; Donkar, 1981). Earnings above the SSA taxable earnings maximum in that first quarter were not recorded, but this was believed to be a modest omission since relatively few workers earned over the annual maximum contribution base more in the first quarter of the year (Myers 1993). For the period 1973-1977, then, the AWI was simply determined by dividing the sum of total taxable earnings reported to the agency during the first quarter of the year, by the number of taxpaying workers (multiplied by four to obtain annual earnings).

1957-1972: For years prior to 1973, it was deemed impossible to reproduce the average computations just described, since the SER first-quarter earnings records were not available. As a result, the SSA turned to a 1-percent sample of annual employee-employer records collected during 1957-1972, a file known as the Longitudinal Employee-Employer Data (LEED). This provided quarterly earnings amounts up to the taxable earnings cap, as well as the number of wage earners in that quarter; consequently, the file permitted the computation of first-quarter taxable average earnings. Recognizing that inconsistencies might result from using different data sources before and after 1973, the agency used LEED earnings only to compute the *annual percentage changes* rather than average earnings *levels*. These 1957-1972 changes were then used to backward-index the AWI level devised previously for 1973. 1951-1956: Before 1957, no LEED data were available, so the agency turned to a 0.1 percent sample of earnings records (up to the taxable cap) collected for 1951-1956. These were again used to derive *changes* in average first-quarter earnings over that period, and these 1951-56 changes were then used to backward-index the AWI from 1957.

1978-present:[8] The 1977 Congressional committee that proposed the AWI concept had intended

that the AWI going forward would rely on 100-percent reporting of all earnings, rather than just

a sample of earners. Accordingly, for 1977 and 1978, the agency collected earnings reported on

personal income tax forms (Form 1040). Thereafter, the committee stipulated that earnings

records would have to come from the Internal Revenue Service (IRS) Wage and Tax Statements,

known as the "W-2 forms." From 1978 on, the SSA contracted with the IRS to derive the data

required to compute the changes needed to update the AWI previously described. Instead of

using the new earnings records to compute new AWI earnings levels, however, the new data

were again used only for computing *changes* in annual earnings, and these were applied to bring

forward the old AWI. Finally, in 1986, the agency began to process W2 data on its own, and

once again these records were used to derive annualized *changes* to update the AWI (Clingman

and Kunkel, 1992).

It is not known whether the AWI series just described, which relies on 30-year old

earnings records, accurately reflects workers' *actual* earnings patterns since the mid-1970s. One

reason to think it does not is clear in Figure 3: the AWI over the entire period appears to have

been about 18 percent higher than the *median* wage and salary income recorded by the US

Census Bureau for all workers.[9] The dollar as well as the percentage gap between the AWI and

median earnings has also grown steadily from about 1980 on. Furthermore, the *mean* earnings of

wage and salary workers depicted in Figure 3 trace a very different course from the AWI: over

the entire period, the AWI came to only 72 percent of national *mean* earnings and the dollar gap

is rising over time. We therefore conclude that the AWI has not been a particularly accurate

measure of mean or median labor earnings from 1951 to 2001: specifically, it is increasingly

---

[8] This section relies on Donkar (1982).
[9] The figure given is the average of male and female workers' wage and salary incomes.

overstating *median* labor earnings patterns, and it consistently has understated *mean* earnings. Below, we discuss whether and how the AWI might be revamped to better track national earnings profiles.

Figure 3 here

Hypothetical Steady Earners: In recent years, analysts have raised new questions about the distributional aspects of social security benefits and taxes within and across cohorts of workers and retirees. This interest, in turn, has prompted the agency to devise simulated or hypothetical profiles often used in policy discussions. One frequently-used profile has been called the medium or "average steady" profile, where a hypothetical individual is assumed to enter the workforce at age 22, remain continuously employed until age 65 (except for periods of disability), and earn the AWI throughout his lifetime (Nichols et al., 2001). Two alternative hypothetical profiles have also been formulated: a "low steady" profile, which is assumed to have the same permanent labor force attachment but earnings worth only 45 percent of the AWI in all years; and a "high steady" profile that also has permanent market attachment but earnings stand at 160 percent of the AWI in all years (Nichols et al., 2001).[10] The resulting profiles appear in Figure 4, in both nominal and constant dollars. It is instructive to note that, while the nominal profiles rose steadily over time, the real profiles were quite volatile, particularly during the high inflation period of the early 1970s.

Figure 4 here

One reason the steady earner concept has been used by the agency is that it relies on the AWI concept familiar to SSA analysts. In addition, it is relatively straightforward to produce: once having calculated AWI, computing steady earner profiles does not require collecting

---

[10] The 45 percent figure may have been developed to resemble the earnings path of a minimum wage worker, but it is unclear why 160 percent was used for the high profile.

additional data on actual work and earnings levels. Nevertheless, the steady earner concept has some shortcomings as a representation of actual workers' pay trajectories. In particular, it ignores life cycle earnings patterns, so that a steady earnings profile of exactly the AWI at all ages will likely overstate actual earnings paid to younger workers, and understate earnings of prime-age workers. Earnings volatility over the life cycle is also ignored, as are fluctuations in pay that could differ across earnings levels. A steady earnings profile also assumes away within-cohort and across-cohort differences in the ages of entering employment, unemployment (or un/nonemployment), and patterns of early retirement. The steady earnings profile also attributes the same percentage boost in pay to all workers irrespective of earnings level, a supposition that may be far from accurate for actual workers. All of these issues imply that AWI profiles and projected benefit amounts inferred for such profiles are likely to be highly inaccurate representations of what actual workers earn and, in turn, receive in the form of benefits. Thus using hypothetical steady profiles to project future retirement benefit amounts could lead analysts to draw incorrect inferences about payroll tax collections and benefit payouts. In addition, policy projections that rely on the steady earner construct may generate inaccurate conclusions regarding the impact of rule changes.

Hypothetical Scaled Earners:[11] As an alternative to the steady AWI earner profile just discussed, the OACT developed three "scaled worker" profiles in 2001 to more accurately reflect actual earning patterns over the lifecycle. In particular, these scaled earnings profiles were structured to allow earnings to rise with age. Nevertheless, the scaled profiles continue to assume that people enter the workforce at age 21 and remain continuously employed until retirement, disability, or death.

---

[11] This section draws heavily on Nichols et al. (2001).

The process of 'scaling' the AWI data involves three somewhat complex steps involving (a) the development of "raw" scaled factors; (b) the implementation of a consistency adjustment; and (c) the determination of annual scaled earnings using "final" scaled factors. In the first step, so-called "raw" scaled factors are developed using social security covered taxable earnings data for currently fully-insured employees who pay payroll tax taken from the agency's Continuous Work History Sample (CWHS). According to agency rules, a fully-insured worker has at least 40 covered employment quarters (or for younger employees, more than six and more than the employee's age minus 21). This datafile was then matched by the agency to administrative records on covered earnings from 1951 forward. Actual earnings figures were then divided by the AWI in each year (and then by 12 months) to produce each individual's Average Indexed Monthly Earnings (AIME). Probably to smooth the data, individual workers' AIMEs were then summed by age over a 10-year period, and divided by the total number of earners of that age. This computation results in so-called *age-specific raw scaled* factors.

Next, these raw scaled factors were adjusted to make them consistent with the low, average, and high AIME amounts for the corresponding steady earner profiles described above. It should be noted that the agency also moderated the raw scaled factors for ages 62-64 because the raw factors generated in the first step exhibited a discontinuous rise for that age bracket; the approach taken was to hold constant nominal earnings from age 62 on. Annual earnings for other ages are then predicted for each hypothetical profile of workers turning age 21 in any given year, by multiplying the raw scaled factor by the AWI for the given year. These predicted profiles are used to adjust the scaled factors so that the AIMEs for low, medium, and high *scaled* profiles are consistent with the AIMEs for low, medium, and high *steady* profiles. This is accomplished by multiplying the adjusted raw scaled factors by the appropriate steady worker

AIME, divided by the AIME derived from the scaled factors. For example, the adjusted scaled

factor for a 35 year old hypothetical low scaled worker is computed using the following formula:

$$(adjusted\ scaled\ factor)_{35,low}=(raw\ scaled\ factor)_{35}\left(\frac{AIME_{steady,low}}{AIME_{scaled}}\right).$$

Finally, the three sets of adjusted scaled factors are then used to produce "final"

hypothetical scaled low, medium, and high earner profiles. This is accomplished by multiplying

the age-specific scale factors by the AWI in each relevant year. The first set of published scaled

factors used earnings data from 1988-97, but these are updated annually using the 10 most recent

years of data. The rationale for using the most recent decade is that recent earnings information

is felt to be the best predictor of future earnings. For instance, the working patterns of women

have changed since the mid-1900s, so using scaled earnings from that time period might not be a

good indicator of their future earnings patterns.

Figure 5 illustrates the scaled profiles for low, medium, and high hypothetical earners

turning age 21 in 1951. Panel (A) shows that both the levels and shapes of the scaled profiles

differ rather dramatically from those in Figure 4. The scaled worker is assumed to have no

earnings growth after age 62 which means that pay profiles are flat from that date forward.

Inflation-corrected earnings, depicted in panel (B), are effectively hump-shaped, with peak

earnings occurring in the mid-40s and declining rather sharply thereafter.

Figure 5 here

A different way to evaluate the impact of the scaling factor is to compare the AWI with

resulting scaled earnings profiles. A graphical representation of this appears in panel (A) of

Figure 6. In nominal terms, the scaled medium hypothetical individual turning age 21 in 1951

would have tracked the AWI fairly closely until about 1987 (in his late 50s). After that, the

scaled profile no longer kept up with inflation and productivity growth. The lower panel shows

that in real terms, the "medium" scaled profile had earnings worth one-third the AWI during the early phase of the worklife, one-quarter higher than the AWI in middle age, and one-third less than the AWI in later life. Similar patterns can be computed for low and high scaled profiles in both current dollar and real terms, and these also diverge notably from the AWI.

Figure 6 here

Compared to the hypothetical *steady* profiles described earlier, the *scaled* profiles have both advantages and disadvantages for analysts seeking to understand the distribution of benefits and payroll taxes. For instance, the scaled pattern takes into account an empirically-based age linked life-cycle earnings profile; for this reason, scaled profiles are more likely to reflect actual workers' earnings patterns. On the other hand, the scaling approach is difficult to understand and explain. Furthermore, the scaling approach does not explicitly recognize periods of non-employment. Years of zero earnings are included in the earnings base via inclusion of zero earnings in the computations, but the method only "counts" the fully insured population and hence it eliminates anyone entering the job market after a long absence (Clingman et al., 2001). Additionally, scaled profiles cover only a 10-year period of data, so they may not accurately represent any cohort's entire lifetime of earnings.

Another concern regarding the hypothetical scaled profiles is that, though the scale factors do vary with age, each cohort is still assumed to have an identical age-earnings life cycle profiles, which is unlikely to be true. To illustrate the point, Figure 7 offers a simple simulation. In Panel A, scaled earnings are used to reflect low pay when young, rising pay in middle age, and a slight fall in pay during the later years. In this example, each of three cohorts is assumed to have an identical age-earnings profile, so the cohorts differ only according to the year in which each turns age 21 (years 1, 6, and 11 respectively). The average of scaled earnings in any

year that includes all three cohorts is represented in Panel A by the solid black line. The scaled age-earnings curves of the two "older" cohorts are higher than the average until about year 30, and they are positioned on or below the average for the remaining years. Conversely, the youngest cohort has lower-than-average scaled factors until year 30 and thereafter has higher-than-average scale factors.

Figure 7 here

What this example illustrates is that a national average of workers' earnings in any given year will almost always misrepresent cohort-specific earnings, even when the cohorts have identical age-earnings patterns. In this case, the average overestimates cohort-specific earnings roughly half the time and understates them the other half. If the goal is to calculate lifetime average earnings so that high and low errors net out at zero, this might not be an issue. But the reader must be reminded that scaled curves in Panel A assume there are only age-specific changes in earnings, without any *cohort*-specific earnings growth. In actuality, of course, earnings growth does differ across cohorts, in which case both the slope and the shape of the profiles will vary across cohorts. Panel B of Figure 7 illustrates such an example, where each cohort is allowed to earn more in real terms than the preceding one.[12] In this simulation, Cohort 2 tracks the average of cohort scaled factors well, but Cohorts 1 and 3 never match the average. Furthermore, the difference in earnings growth rates produce average cohort scaled factors that always underestimate earnings for Cohort 1 and overestimate earnings for Cohort 3. In Panel C we graph the difference between scaled factors for each cohort from Panel B and the average.

The example is offered to illustrate the potential for error when estimating earnings using year-specific as opposed to cohort-specific scaled age-earnings factors. For actual workers, these

---

[12] Panel B assumes that Cohort 1 experiences an additional real earnings growth rate of 1.8% per year above the scaled factors; Cohort 2 earns an additional 1.3%; and Cohort 3 an additional 0.98%; these are simply illustrative figures and not necessarily representative of any specific birth group.

differences could be more or less important, so in the next section we compare actual earnings profiles with the stylized profiles developed by the SSA.

Comparing Hypothetical with Actual Workers

We turn next to an analysis of actual workers' lifetime earnings profiles, for which we rely on the nationally representative Health and Retirement Study (HRS).

Design and Methods: The HRS, along with associated linked earnings records from SSA, provides a unique opportunity to assess lifetime earnings patterns.[13] As described in our earlier work (Mitchell et al. 2003), the initial HRS cohort was age 51-61 when first interviewed in 1992 (and spouses of any age were also interviewed). In addition to containing rich health and demographic information, the linked HRS datafile provides a comprehensive picture of workers' lifetime earnings patterns. [14]

In what follows, we use the covered earnings and W-2 earnings histories to compare actual workers' *covered and total* earnings paths with AWI calculations as well as the low/medium/high hypothetical profiles described above. For the period 1950-91, earnings up to the taxable maximum were available in nominal terms, and from 1980 to 1991, "W-2" earnings were also available which include labor compensation above the taxable earnings ceiling.

Findings on Employment Patterns: As noted above, the scaled and steady hypothetical profiles do not explicitly recognize periods of unemployment or zero earnings. Whether this

---

[13] For an extensive discussion of the Health and Retirement Study and the data that may be linked with respondent survey information under restricted access conditions, see Mitchell et al. (2000).

[14] Data were obtained for a majority of HRS respondents and spouses providing permission to link their survey data with administrative records supplied by the Social Security Administration and also with pension plan descriptions provided by respondents' employers; see Mitchell et al. (2000). A match with SS earnings records was feasible for approximately 75% of the respondents; in addition we dropped approximately 700 married respondents from the analysis due to missing data on spouse lifetime earnings. While omitting these non-matched cases might bias the sample, if those who with a matched file differ from those lacking a match, our analysis indicates little reason for concern.

assumption squares with real workers' patterns can be evaluated using Figure 8, Panel A. In the early portion of 1950's (when respondents were in their early 20's), young men had rather high rates of non-employment, with as many as 46 percent of HRS respondents indicating no positive earnings in any given year. By contrast, only about 5-7 percent of the sample of men had no positive earnings reported in any year after their 20's. For HRS women, the picture is even more striking: typically more than 40 percent of women in this age group had no positive earnings from 1951 to 1971 during the peak childbearing and rearing period. By the 1980's, when the cohort reached its 50s, rates of zero earnings still were high, at more than 20 percent. For both sexes, then, but particularly for women, the high rate of zero earnings observed suggests that the hypothetical earnings profiles with no zero-earnings years will generate far higher computed lifetime earnings than 'actual' workers in this cohort.[15]

Figure 8 here

Findings on Earnings Patterns: We also indicate how often annual administrative earnings records indicated that workers were at the taxable earnings maximum (Panel B, Figure 8). This is important since the HRS data reports total earnings from 1980 on, but only taxable earnings prior to that date. For HRS women, the capping was not much of a concern: their earnings attained the taxable cap an average of only 2.3 percent of the time between the years 1951 and 1979 when no W-2 records were available. Among HRS men, however, the prevalence of capped earnings was higher, averaging about 15 percent during the 1950's, 44 percent in the

---

[15] We omitted from the HRS sample workers who had at least two years of zero taxable earnings in the 1980-91 period but who also had positive W-2 earnings. This curtails the chances that HRS respondents used in the analysis might have been working in government or other non-covered sectors. Our baseline sample also excludes respondents who have zero total lifetime earnings over the period measured by the data.

1960's, and 47 percent in the 1970's. For years prior to 1980, we correct for capping by

generating predicted values of year-specific real earnings for each capped worker.[16]

Since the AWI averages earnings across all covered workers in any given year, the

approach might not effectively represent earnings of people in any given birth cohort. To assess

this point, Figure 9 compares the AWI pattern to earnings of HRS respondents arrayed by birth

cohort; there are 11 birth cohorts in our datafile, representing people born 1931-1941. Panel A

compares *mean* annual earnings by birth cohort, to the AWI. Here we see that the age-earnings

profiles for each cohort's average earnings follow similar patterns, first rising and eventually

flattening out, with a few relatively large values producing spikes in average earnings during the

1970's. It is of interest that the AWI profile is higher that the average earnings for all birth

cohorts for about two decades, from 1951 to 1971. But over 1970-1991, most of the cohorts

exhibit higher average earnings than the AWI series. Evidently, given the shape and position of

each cohort's earnings relative to the AWI, it is unlikely that the AWI would effectively

represent average earnings of any of the eleven cohorts.

Figure 9 here

We replicate the same data in Panel B but mitigate the influence of possible outliers using

the median 10% of annual actual earnings; the median 10% computes an average from the

middle 10% of the earnings distribution in each year. Clearly this smoothes actual earnings and

also lowers the curve below the AWI series. The lesson appears to be that the AWI does not

---

[16] In particular we estimated right-censored Tobit models with panel-level random-effects of the natural log of annual earnings, with explanatory variables including age, education, race, sex, parents' educational attainment, and year effects; differential effects were allowed by marital status. Predicted earnings were then substituted for capped earnings if they were greater than the taxable maximum in years prior to 1980. Results suggest (cf Appendix Table 1) that women earned less than men, whites earned more than nonwhites, and earnings rose with age and education. Our approach is similar to those implemented by Scholz et. al (2003) and Engelhart and Cunningham (2002).

account for workers dropping out of the labor force, so it systematically tends to overstate lifetime earnings.

While the AWI does not track cohort-specific earnings well, it may be that profiles constructed using the AWI and scaled factors do better in representing "typical" workers. Unfortunately, they do not. This is evident from Figure 10 which compares the low, medium, and high scaled profiles of the median 10% of earnings for all HRS workers in the 1936 birth cohort, [17] as well as those who have long enough work histories to be insured for retired worker benefits from social security. [18] The *high scaled* profile is substantially above the 1936 HRS birth cohort, and the *medium scaled* profile, which some might think should represent the "typical" worker, was in fact above the median 10% of earnings in all years. It is interesting that the *low scaled* profile tracks the HRS median 10% until the late 1960's, when the series began to diverge. From then on, the HRS median 10% rose, relative to the low scaled profile. It is true that the medium scaled profile better tracks the fully-insured workers' median earnings rather than all taxpayers.

Figure 10 here

To make the more general point, we next compute the distribution of actual lifetime earnings for all eleven birth cohorts in the HRS sample (in $1992). For comparative purposes, we also compute lifetime earnings for scaled hypothetical workers assuming that they had been born in each of the same years (1931-1941), and evaluate the weighted average of these profiles,

---

[17] Results for other HRS birth cohorts are similar.

[18] Thus far, the HRS sample examined included only respondents with some positive earnings over the period of observation (1951-1991). Some of these persons did not contribute long enough to be insured for retired worker benefits from social security., and one might think that scaled profiles were intended to look like "typical" social security beneficiaries, rather than all contributing workers. To this end we also compute earnings trajectories for workers who have accumulated 40 "Quarters of Coverage" so as to be insured for Social Security retirement benefits at age 62. Workers earn up to four quarters of coverage a year based on the level of payroll tax contributions in a given year. For example, in 2002 a worker would earn one quarter of coverage for earnings of $870, and could earn the maximum of four quarters in a year with covered earnings of $3,480 or more. For more information of insured status see Mitchell et. al (2001).

assuming they were represented in the same proportion as in the HRS sample. This permits a comparison of the weighted average of the scaled profile with the distribution of average lifetime earnings for our actual HRS sample.

Figure 11, Panel A, offers the comparison for the HRS baseline sample, while Panel B reports the same results but focuses only on the fully insured sample. Overall, it is clear that earnings implied by the scaled profiles are substantially higher than comparative values using actual HRS earnings data. For example, the implied medium *scaled* profile of $20,185 lies at about the 67[th] percentile of average *actual* lifetime HRS earnings and it is more than 50% above the lifetime earnings of the median HRS worker. Further, the medium scaled profile is roughly worth 150 percent of the HRS actual medium average earnings level ($13,313). The low scaled profile ($9,058) also looks high, at around the 42[nd] percentile of the actual HRS distribution.[19] Lifetime earnings for the high scaled profile stand at the 82[nd] percentile of the actual HRS distribution ($32,302), rather than the approximately $20,000 in average lifetime earnings that the actual data would predict (between the 65[th] and 70[th] percentiles). In other words, all of the scaled earnings profiles are substantially higher than their target values in the HRS sample.

Figure 11 here

Restricting the sample to fully-insured workers reveals that scaled profiles are still high: that is, the medium scaled profile is about 28% higher ($4,400) than the HRS median, and the low scaled profile is 27% above (about $2,000) the HRS median. Over the course of a working career, this earnings difference is rather substantial, with the gap between the medium scaled profile and the HRS fully insured median worker over 35 years totaling about $150,000 ($1992).

---

[19] The low scaled profile was intended to represent earnings worth 45 percent of the medium earner—in the HRS baseline sample, this would be worth about $5,500 which would be between the 25[th] to 30[th] percentile.

Discussion and Conclusions

This project compares actual lifetime earnings with hypothetical steady and scaled earnings profiles. The hypothetical profiles are often used for distributional analysis despite the fact that the concept on which they are based was not developed for this purpose. In principle, the hypothetical profiles do not reflect any specific cohort's experience, either in terms of their earnings when working, or actual labor force attachment profiles. In practice, when we compare the hypothetical profiles with actual earnings, we find that the scaled profiles do not compare well to actual earnings paths.

One reason for this divergence is that the assumption of steady work does not track the experience of actual HRS workers. Well over one-third of all men and women in our sample did not have covered earnings in their 20s, and many women had zero earning years after that. All the hypothetical profiles are higher and flatter than the typical HRS workers in our sample. We also find that the Average Wage Index, intended to reflect a weighted average of actual earnings at any given time, does not match the average earnings of any given cohort. In addition, the AWI exceeds average actual earnings during working cohorts' early years, and, using measures unaffected by high outlier earners, it is still higher than HRS actual cohort earnings in all years. Further, median HRS actual earnings were more similar to the low versus the medium scaled profile. Even after restricting the HRS sample to respondents with substantial work histories, the medium scaled profile is 28% above HRS actual median earnings, implying a lifetime difference of more than $150,000 ($1992).

An implication of this finding is that benefits received and taxes paid for a hypothetical medium scaled profile would substantially misstate benefits received and taxes paid by a typical HRS worker. One would also anticipate that benefit replacement rates for hypothetical profiles

would also be much lower than actual replacement rates for actual workers. Though the scaled profiles do better at representing fully insured workers, they are still not representative of the actual distribution of earnings histories of all who pay payroll taxes to the system.

Another implication of our research is that policy simulations of possible social security reform proposals should use earnings data representative of workers in the real-world labor market. For example, recent proposals to modernize the social security system have emphasized the need to enhance retirement security for those at the bottom of the lifetime earnings distribution (Cogan and Mitchell, 2003). To illustrate that actual profiles matter in assessing such a proposal, we simulate the potential accrual of benefits that might have built up in a personal account-type program had one been in place for workers entering the labor market in the mid-1950s; see Figure 12. We use actual earnings data for the 1936 HRS birth cohort and assume 2 percent of pay was contributed annually and invested in assets paying a 4 percent real return.[20] The hypothetical worker on the much higher AWI scaled profile would be projected to accrue twice as much wealth instead of the actual HRS earner.

Figure 12 here

In sum, scaled profiles were produced in the past to meet objectives other than those they are now being asked to meet. They were devised when technology and recordkeeping limitations dictated limited reliance on actual earnings histories. Today, however, new technology and more accurate recordkeeping offer the opportunity to better measure and model actual workers' lifetime earnings for policy and distributional purposes. Analysis that relies on scaled profiles may substantially misrepresent real-world benefit payouts and tax revenue.

---

[20] The selection of specific contribution rates and rates of return is arbitrary and does not affect the pattern of results.

References

Advisory Council on Social Security. 1997. *Report of the 1994-1996 Advisory Council on Social Security, Volume I: Findings and Recommendations*, Social Security Administration, Washington, DC.

Bosworth, Barry, Gary Burtless, and Eugene Steuerle. 1999. "Lifetime Earnings Pattern, the Distribution of Future Social Security Benefits, and the Impact of Pension Reform," CRR WP 1999-06, Center for Retirement Research at Boston College, Chestnut Hill, MA, December.

Butrica, Barbara A., Howard M. Iams, and Karen E. Smith. 2003. "It's All Relative: Understanding the Retirement Prospects of Baby Boomers," CRR WP 2003-21, Center for Retirement Research at Boston College, Chestnut Hill, MA, November.

Clingman, Michael and Jeffery Kunkel. 1992. "Average Wages for 1985-90 for Indexing Under the Social Security Act." Actuarial Note 133, Office of the Chief Actuary, Social Security Administration, Baltimore, Md.

Cogan, John F. and Olivia S. Mitchell. 2003. "Perspectives from the President's Commission on Social Security Reform." *Journal of Economic Perspectives*. 17(2). Spring.

Commission to Strengthen Social Security. 2001. *Strengthening Social Security and Creating Personal Wealth for All Americans, The Final Report*. Washington, D.C. December. www.csss.gov.

Donkar, Eli N. 1981. "Average Wages for Indexing Under the Social Security Act and the Automatic Determinations for 1979-81." Actuarial Note 103, Office of the Chief Actuary, Social Security Administration, Baltimore, Md.

Donkar, Eli N. 1982. "Average Wages for 1980 for Indexing Under the Social Security Act and the Automatic Determinations for 1982." Actuarial Note 112, Office of the Chief Actuary, Social Security Administration, Baltimore, Md.

Engelhardt, Gary and Christopher R. Cunningham. 2002. "Federal Tax Policy, Employer Matching, and 401(k) Saving: Evidence from HRS W-2 Records" NTJ, 55:3 (September): 617-645.

Gustman, Alan L., and Thomas L. Steinmeier. 2001. "How Effective is Redistribution Under the Social Security Benefit Formula?" *Journal of Public Economics* 82(1): 1-28.

Mitchell, Olivia S., Brett Hammond, and Anna Rappaport. 2000. *Forecasting Retirement Needs and Retirement Wealth*. Pension Research Council. Philadelphia, PA: University of Pennsylvania Press.

Mitchell, Olivia S., John W. R. Phillips, and Andrew Au. 2003. "Retirement Wealth and Lifetime Earnings Variability." Pension Research Council Working Paper, Wharton School, July. http://rider.wharton.upenn.edu/~prc/PRC/WP/WP2003-4.pdf

Mitchell, Olivia S. and John W. R. Phillips. 2001. "Eligibility for Social Security Disability Insurance." University of Michigan Retirement Research Center Working Paper 2001-011, University of Michigan, Ann Arbor, MI. June. http://www.mrrc.isr.umich.edu/content.cfm?section=research&content=research_detail&pid=UM00-06.

Myers, Robert. 1993. *Social Security*. 4th ed. Philadelphia: University of Pennsylvania Press.

Nichols, Orlo, Michael Clingman, and Milton Glanz. 2001. "Internal Real Rates of Return Under the OASDI Program for Hypothetical Workers." Actuarial Note 144, Office of the Chief Actuary, Social Security Administration, Baltimore, Md.

O'Harra, Josh, John Sabelhaus, and Michael Simpson. 2004. "An Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model." CBO Technical Paper 2004-01, CBO, Washington, DC.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2003. "Are Americans Saving 'Optimally' for Retirement?" Economics Department Working Paper, University of Wisconsin – Madison, Madison, WI. December, 2003.

Social Security Administration. 2002. *Annual Statistical Supplement to the Social Security Bulletin*. Social Security Administration, 500 E Street, SW, Washington, DC, 20254. December.

Technical Panel on Assumptions and Methods. 2003. *Report to the Social Security Advisory Board*. Washington, D.C. October.

Figure 1: Comparisons of the Consumer Price Index (CPI-U) and the Average Wage Index (AWI): 1951-2001

(A) Comparison of AWI and CPI Changes: 1951-2001
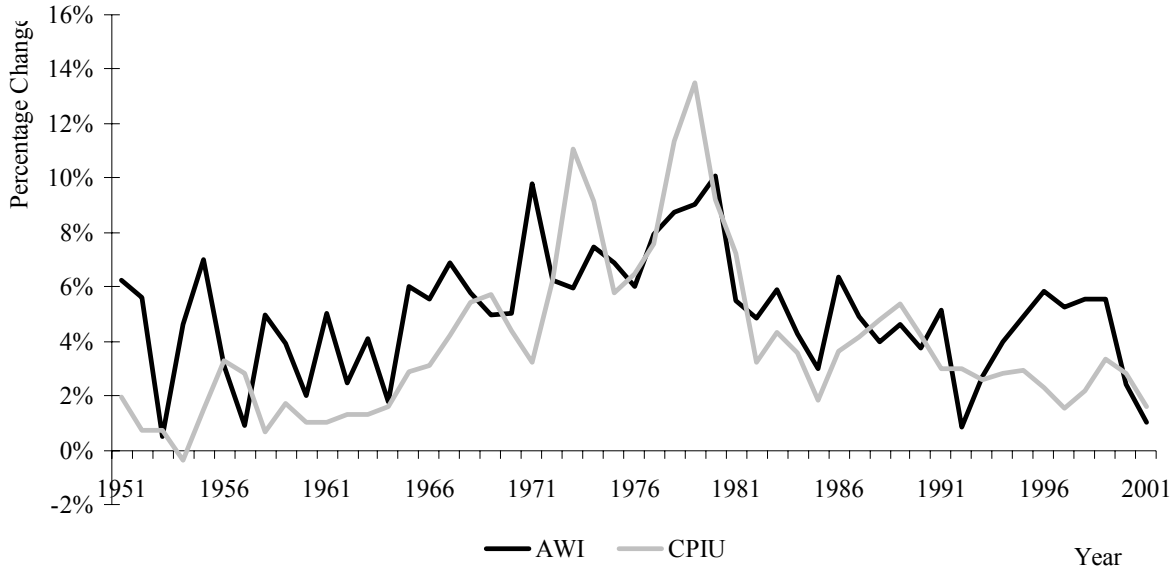
**(B) AWI and CPI-U Annual Adjustment Factors for Worker Attaining Age 60 in 2001**
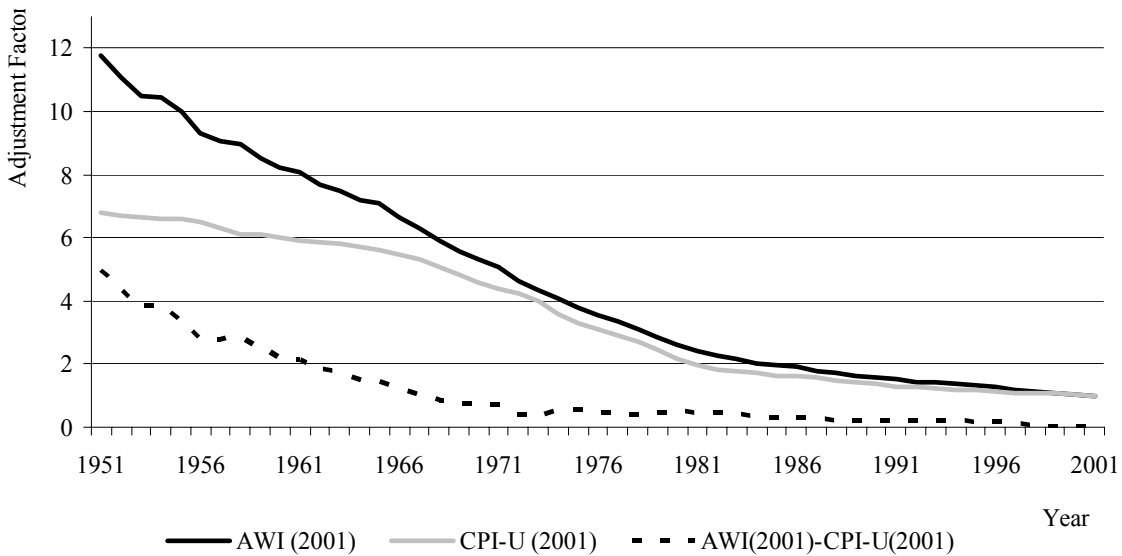
Source: AWI data from SSA website (http://www.ssa.gov/OACT/COLA/AWIgrowth.html) and CPI-U data from the Bureau of Labor Statistics website (http://www.bls.gov/).

Figure 2: A Short History of the Average Wage Index Computations

Period	Average Wage Index Determination
1951-57	0.1-percent sample of first quarter average taxable wages used to index backwards from 1957 figure.
1957-73	Longitudinal Employee-Employer Data (LEED), a 1-percent sample of annual employee and employer records used to index earnings backwards from the 1973 figure.
1973-77	First quarter wage data from SSA Summary Earnings Records (SER) data used to produce average wage.
1977-85	IRS provided 1040 & W2 of complete earnings data to create figures used to index earnings forward from 1977.

Source: Donkar (1981), Donkar (1982), and Clingman et. al (1992). See references for details.

**Figure 3: Comparing AWI and Census Wage and Salary Incomes: 1951-2001**
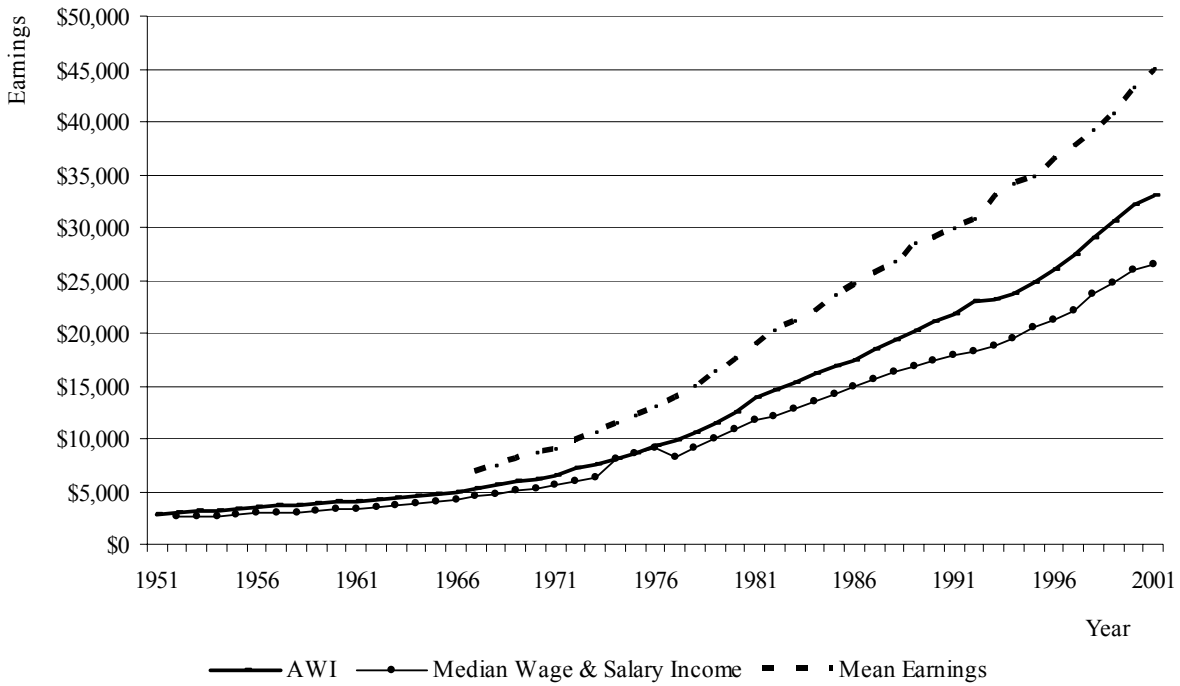
Source: AWI data from SSA website (http://www.ssa.gov/OACT/COLA/AWIgrowth.html), Census Median Wage and Salary Income data and mean earnings data come from the Census website (http://www.census.gov/hhes/income/histinc/p53.html and http://www.census.gov/hhes/income/histinc/p39.html, respectively).

Figure 4: Hypothetical Steady Earnings Profiles: 1951-2001

(A) Hypothetical Steady Profiles: Nominal
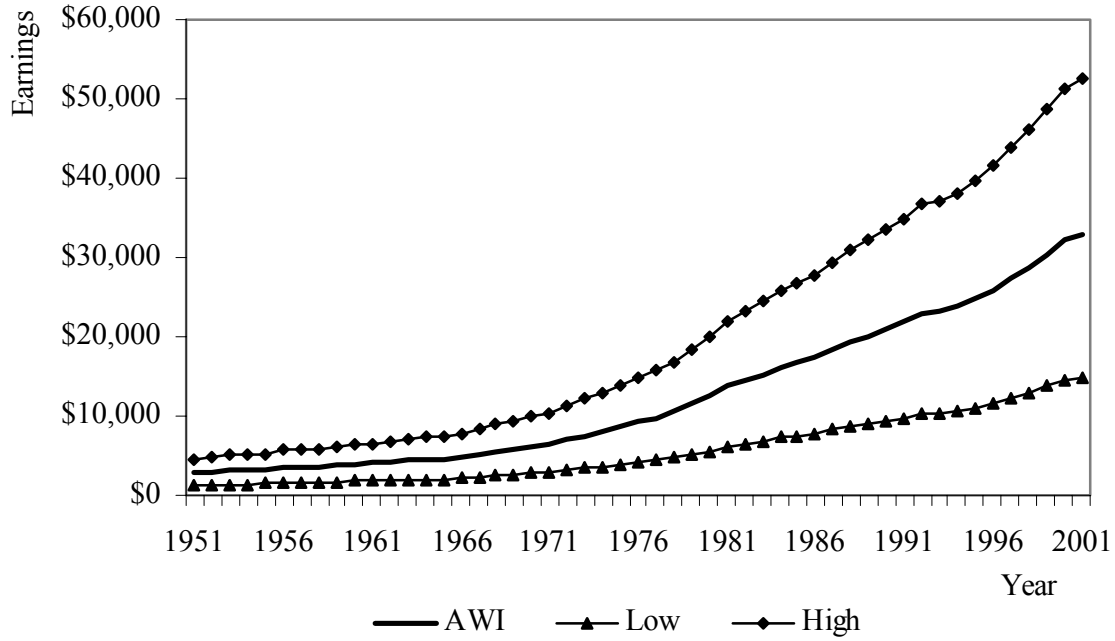
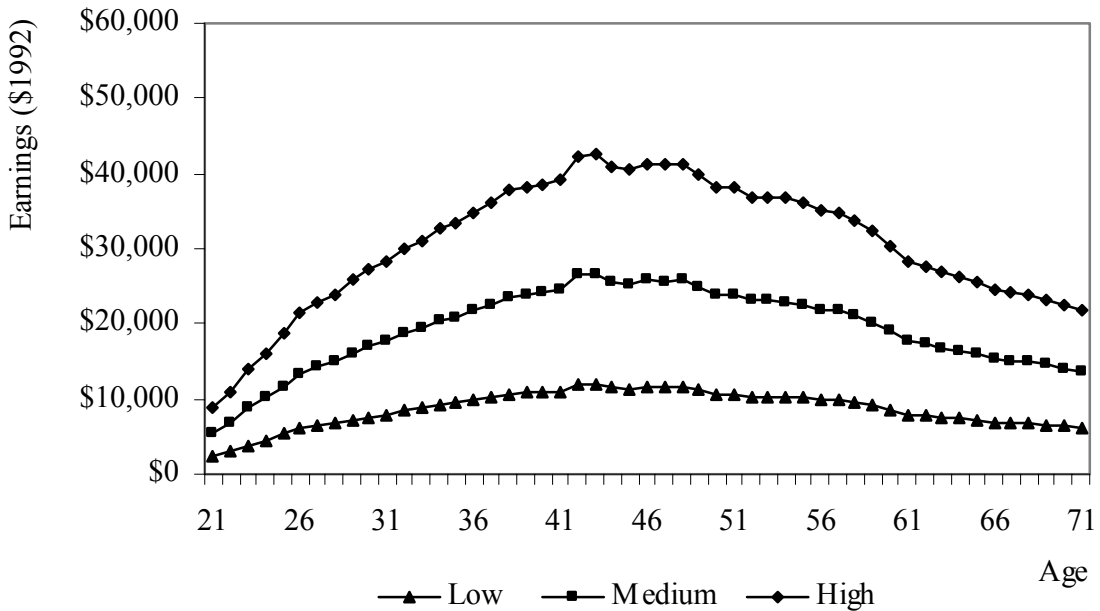**(B) Hypothetical Steady Profiles: Real**
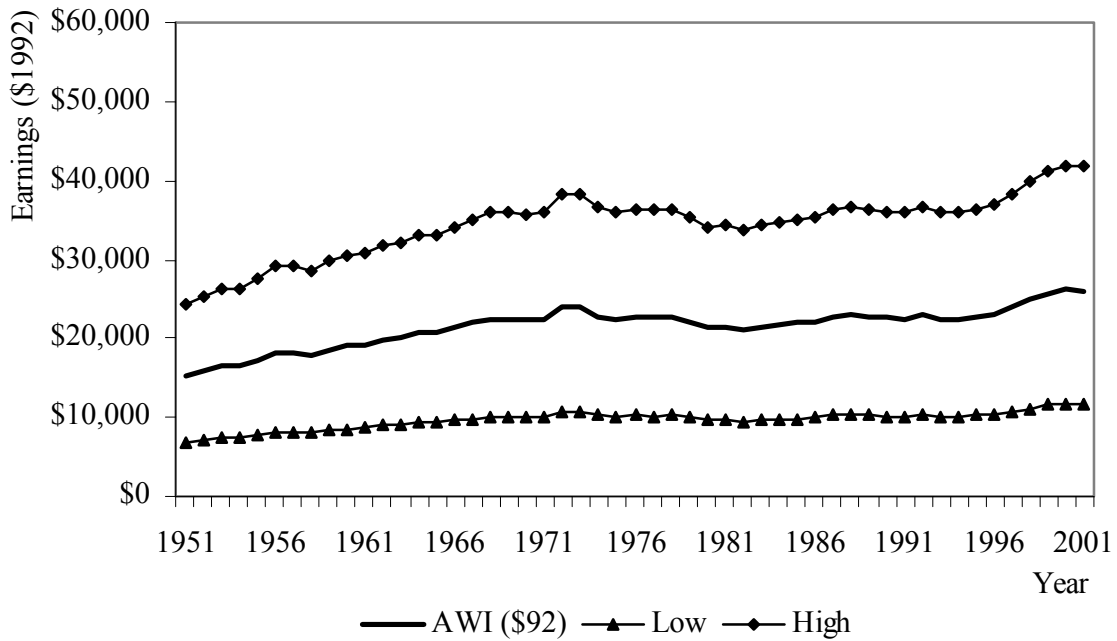
Source: Authors' calculations based on data and methodology from Nichols et. al (2001).

Figure 5: Scaled Steady Earnings Profiles: Age 21 in 1951

(A) Scaled Profiles for Hypothetical Worker age 21 in 1951: Nominal
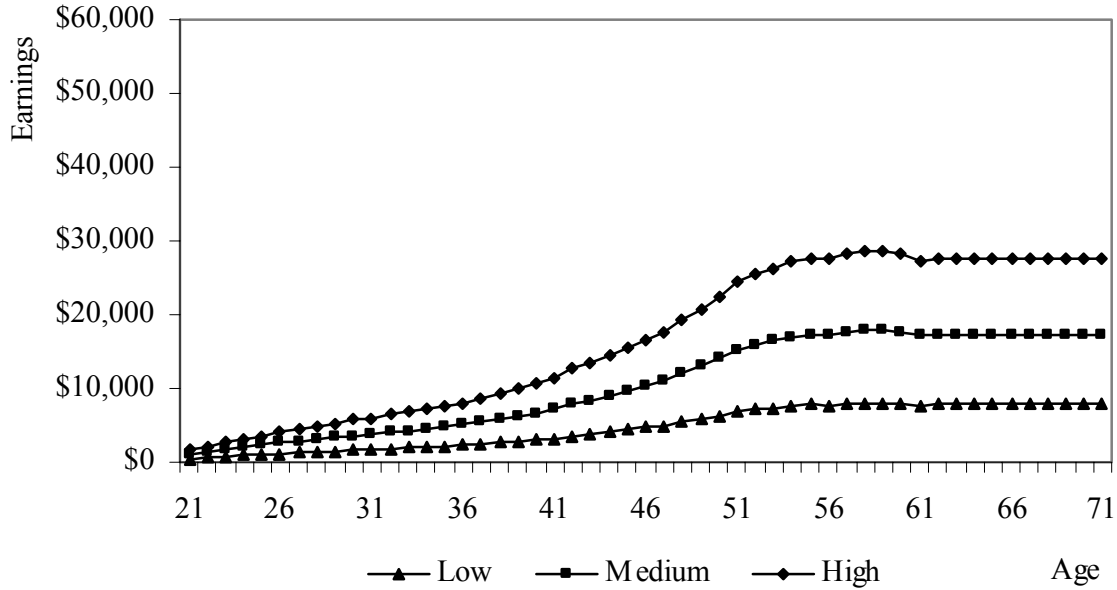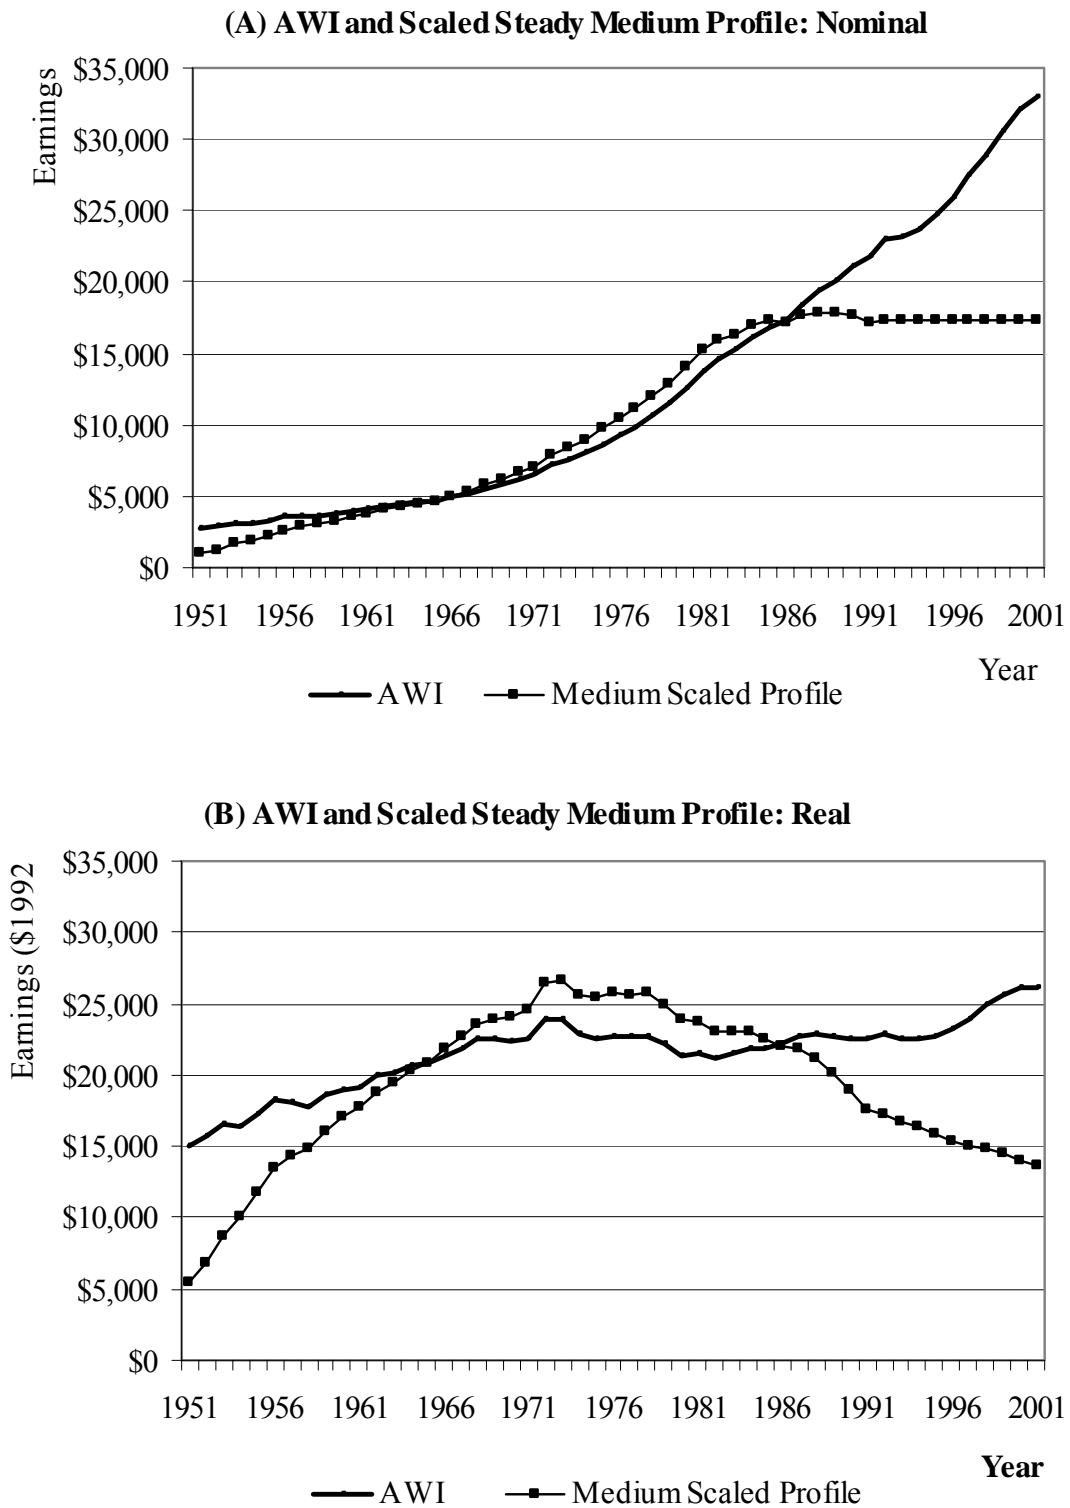

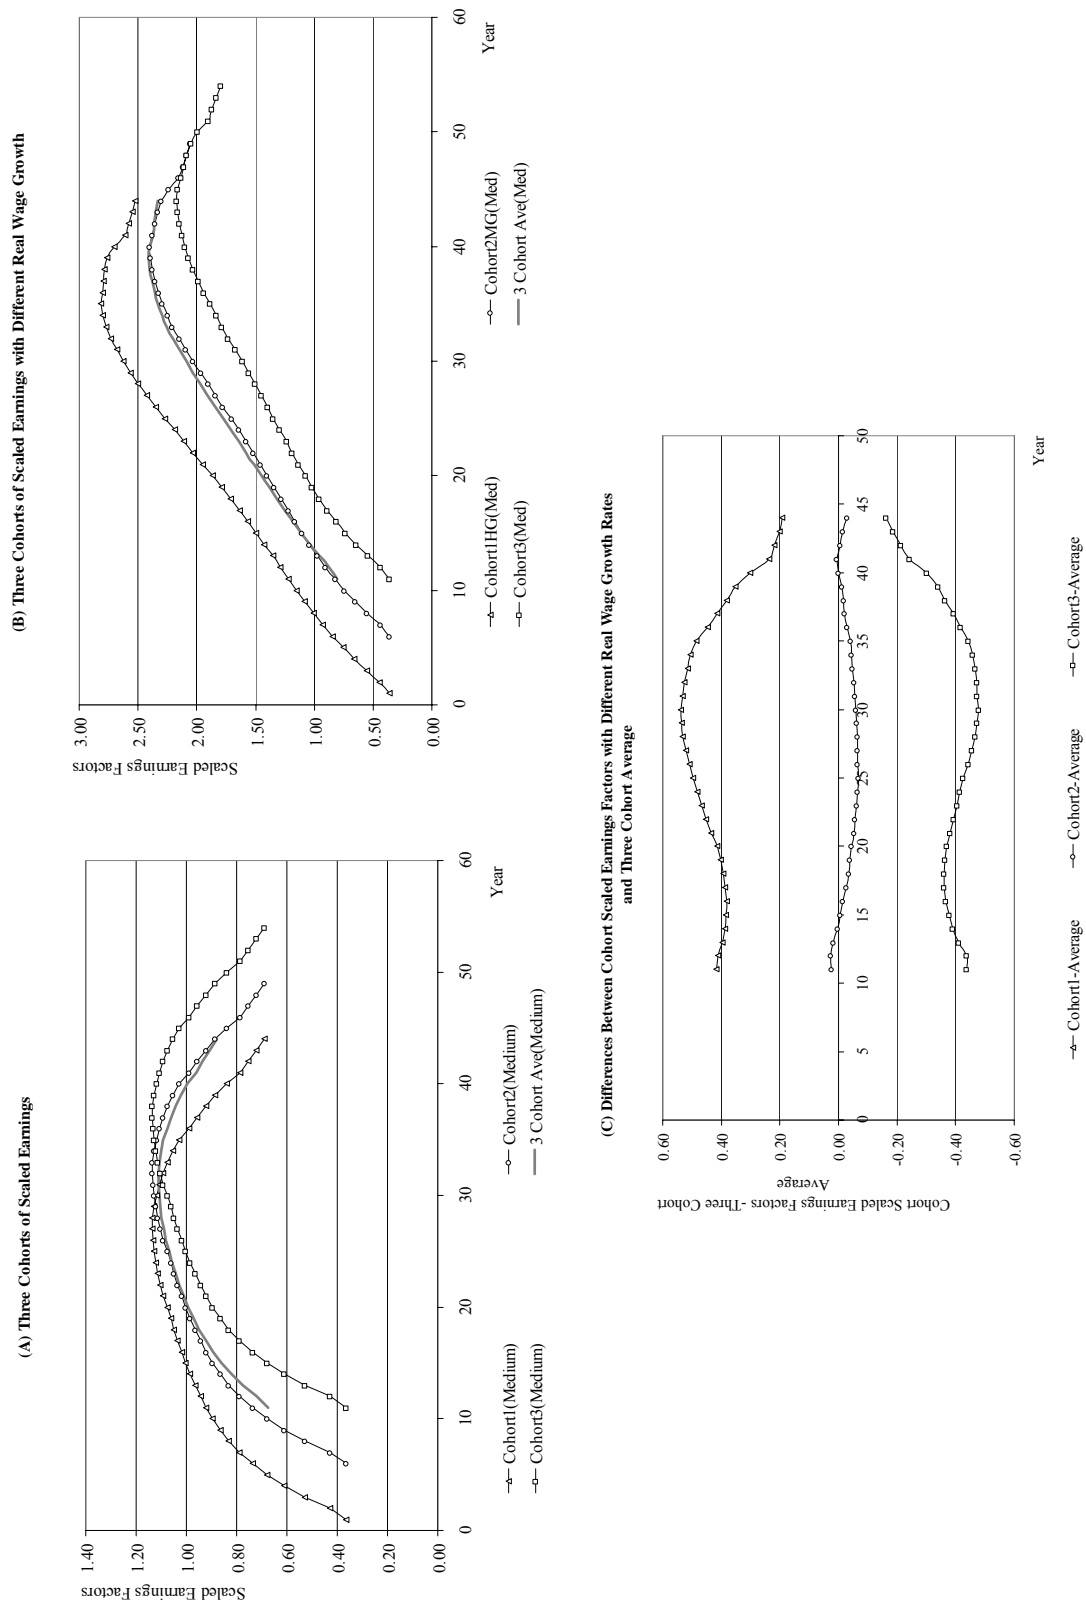**(B) Scaled Profiles for Hypothetical Worker Age 21 in 1951: Real**

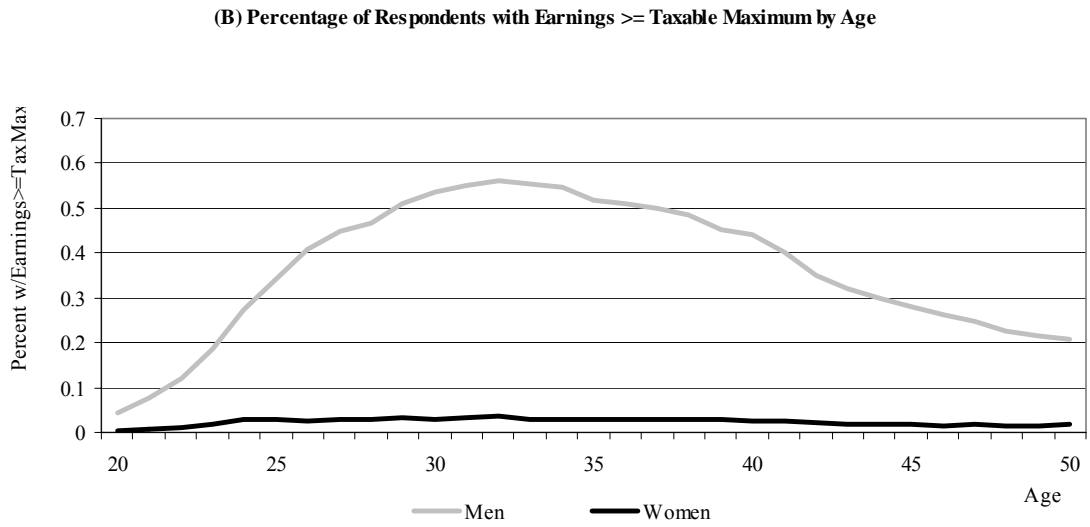Source: Authors' calculations based on data and methodology from Nichols et. al (2001).

Figure 6: **Comparing the Scaled Steady Profile to the AWI: Worker Attaining Age 21 in 1951**

(A) AWI and Scaled Steady Medium Profile: Nominal

(B) AWI and Scaled Steady Medium Profile: Real

Source: Authors' calculations based on data and methodology from Nichols et. al (2001).

Figure 7: Comparing Three Cohorts of Synthetic Earnings Factors to the Cohort Average

(A) Three Cohorts of Scaled Earnings

(B) Three Cohorts of Scaled Earnings with Different Real Wage Growth

(C) Differences Between Cohort Scaled Earnings Factors with Different Real Wage Growth Rates and Three Cohort Average

Source: Authors' calculations using Medium Scaled Factors from Nichols et al. (2001).

29

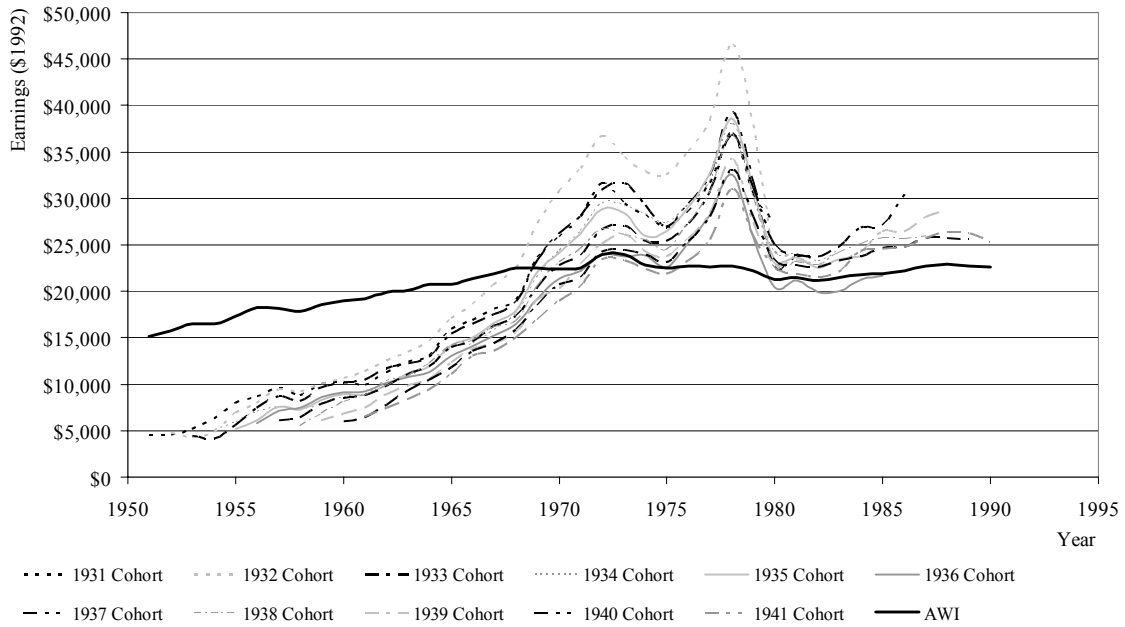**Figure 8: Incidence of Zero and Capped Annual Earnings in the HRS Sample Born 1931-1941**

(A) Percentage of Respondents with Zero Earnings by Age

(B) Percentage of Respondents with Earnings >= Taxable Maximum by Age
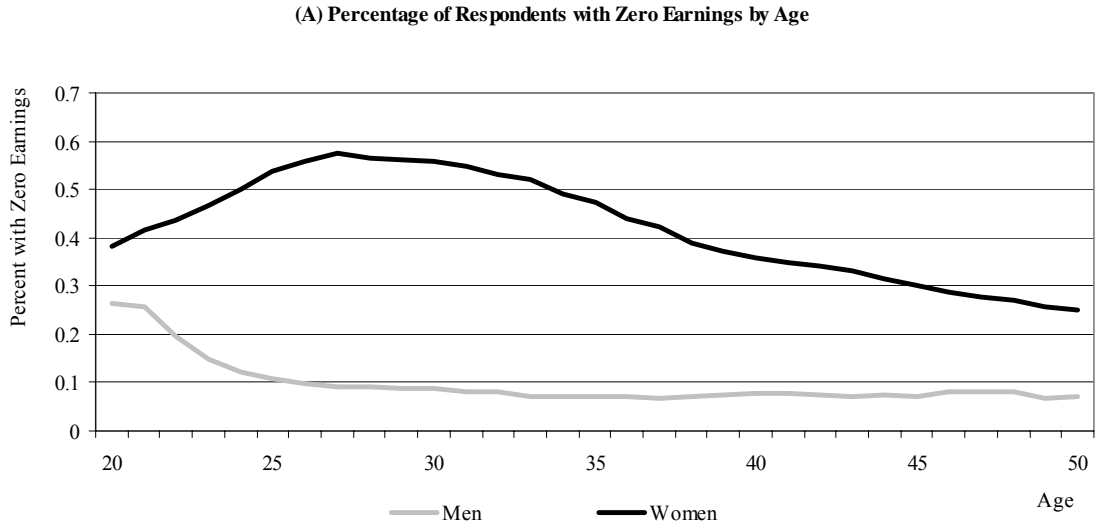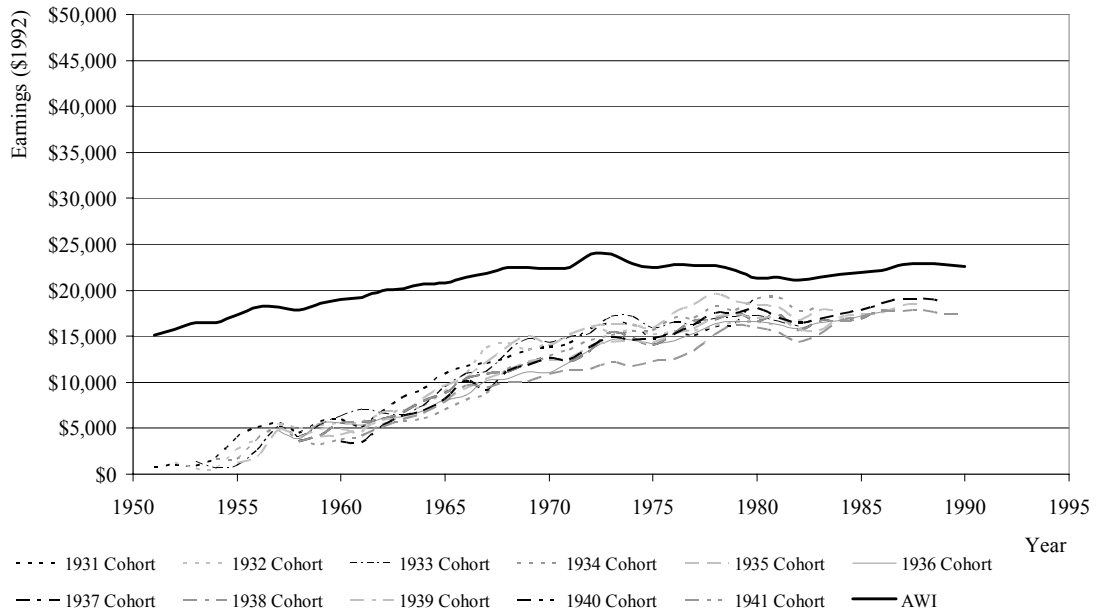

Source: Authors' calculations using the Health and Retirement Study matched to SSA administrative earnings records.

Figure 9: Comparison of AWI and HRS Birth Cohort Earnings: 1951-1991

**(A) A Comparison of AWI and Average Earnings: 1931-1941 Cohorts from 1951 to 1991**

· · · · 1931 Cohort   · · · · 1932 Cohort   — · — 1933 Cohort   · · · · · · 1934 Cohort   —— 1935 Cohort   —— 1936 Cohort
— · · — 1937 Cohort   — · — 1938 Cohort   — — 1939 Cohort   — · — 1940 Cohort   — — 1941 Cohort   —— AWI

(B) A Comparison of AWI and the Median 10% of Earnings:
1931-1941 Cohorts from 1951 to 1991

· · · · 1931 Cohort   · · · · 1932 Cohort   — · · — 1933 Cohort   · · · · 1934 Cohort   — — 1935 Cohort   —— 1936 Cohort
— · — 1937 Cohort   — — 1938 Cohort   — — 1939 Cohort   — · — 1940 Cohort   — · — 1941 Cohort   —— AWI

Source: Authors' calculations using HRS sample matched to SSA administrative earnings and AWI data from SSA (2002).

Figure 10: Comparing Median 10% Earnings with Scaled Worker Earnings: Entire 1936 Cohort and Fully Insured Subset ($1992)
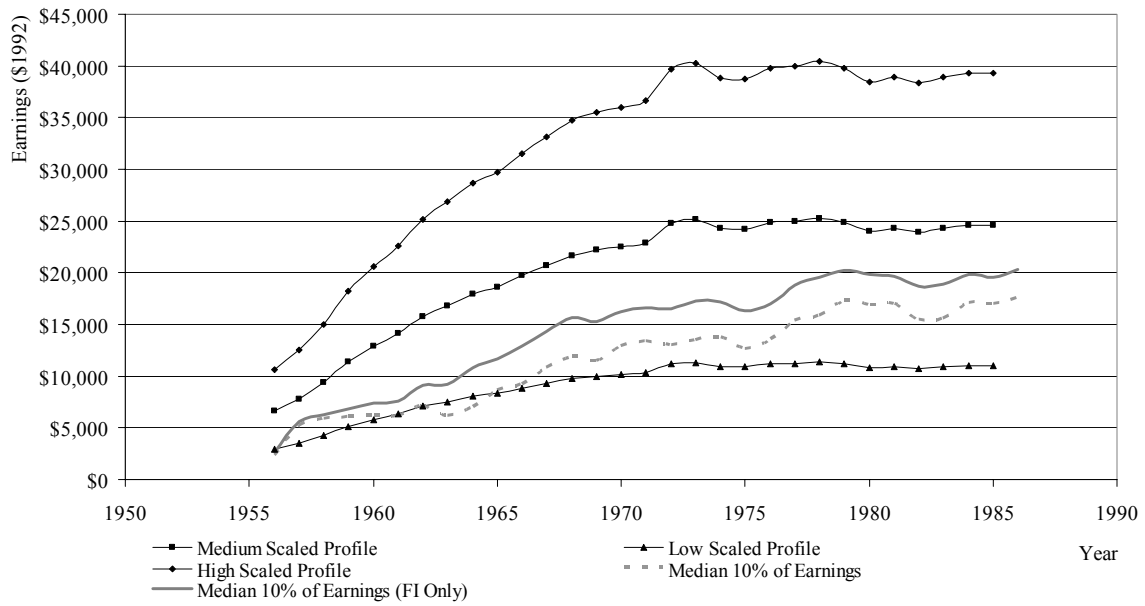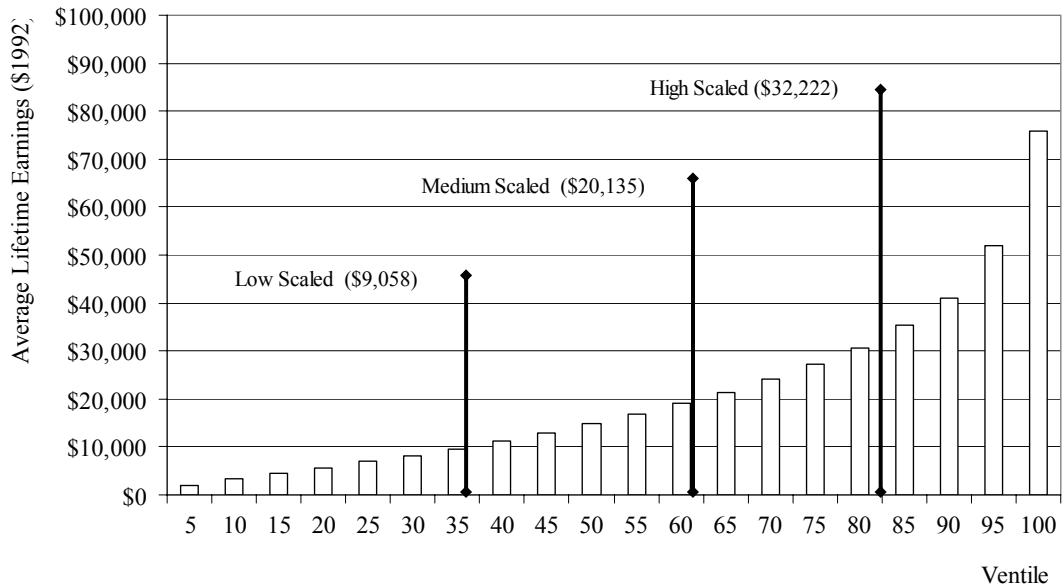

Source: Authors' calculations using (1) a sample of HRS respondents born in 1936 matched to SSA administrative earnings data and (2) scaled worker earnings using low, medium, and high scaled factors and the AWI to generate a scaled worker born in 1936. See Nichols et al. (2001) for details on calculating scaled worker earnings.

Figure 11: Comparing Scaled Profiles to HRS Actual Lifetime Earnings Distributions

**(A) Scaled Profiles Compared to HRS Average Lifetime Earners**
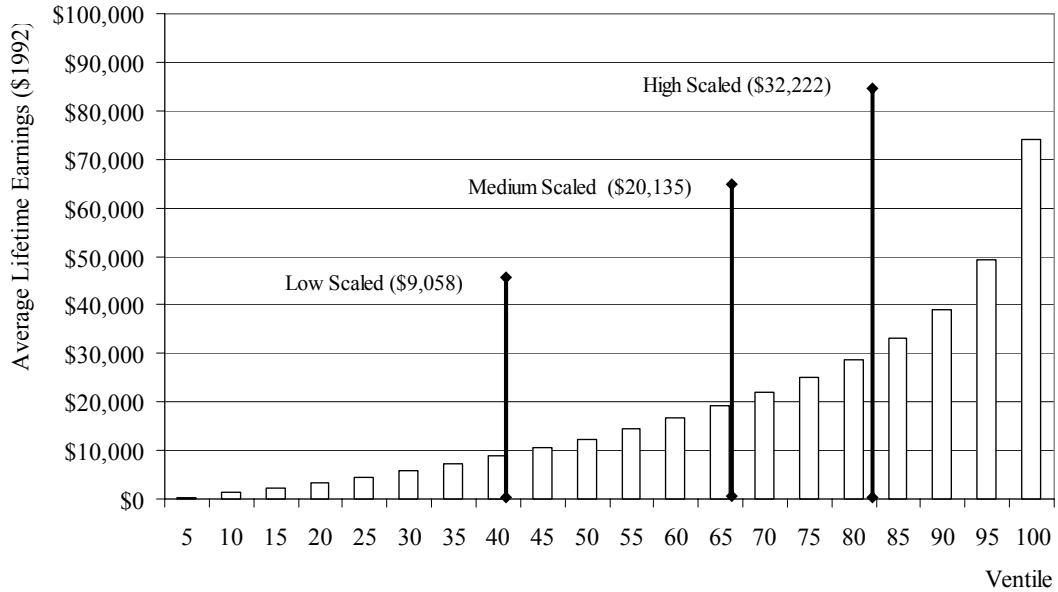
(B) Scaled Profiles Compared to Fully Insured HRS Average Lifetime Earners

Source: Authors' calculations using HRS respondents matched to SSA administrative earnings data and scaled worker earnings using low, medium, and high scaled factors and the AWI to generate average lifetime earnings for the weighted average of scaled workers born from 1931-1941. The weights match the cohort weight observed in the HRS. See Nichols et al. (2001) for details on calculating scaled worker earnings.

Figure 12: Accumulations in a Personal Retirement Account Under Alternative Lifetime Earnings Patterns: 1936 Cohort
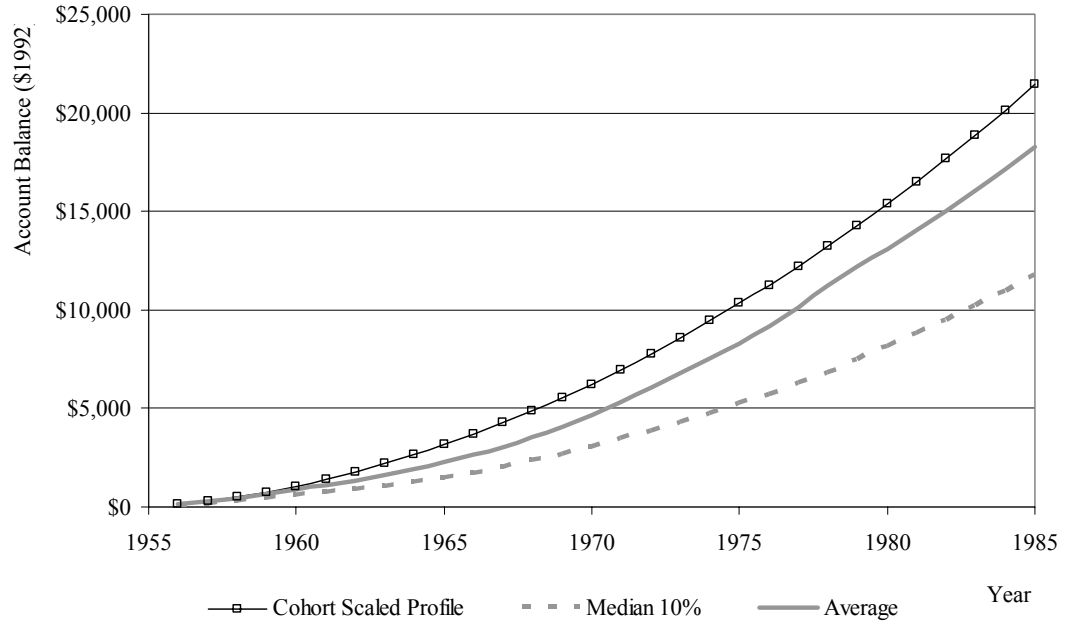
Source: Authors' calculations using a sample of HRS respondents born in 1936 matched to SSA administrative earnings data for the "Average" and "Median 10%", and medium scaled factors and the AWI to generate a scaled profile for the 1936 birth cohort. See Nichols et al. (2001) for details on calculating scaled worker earnings. We assume 2 percent annual contributions to the PRA and a 4 percent real rate of return.

Appendix Table: Tobit Earnings Regression[21]

Variable	Coefficient (SE)	
	Married Sample	Unmarried Sample
Female	-5.78*	-3.39*
	(0.001)	(0.002)
White	1.54*	0.40*
	(0.002)	(0.003)
Age	0.05*	0.09*
	(0.00)	(0.000)
Education (yrs)	0.14*	0.18*
	(0.000)	(0.000)
Mother's educ. (yrs)	0.08*	0.04*
	(0.000)	(0.001)
Father educ. (yrs)	-0.04*	0.02*
	(0.000)	(0.000)
Variance(v_i)	2.44	2.34
	(0.001)	(0.001)
Variance(ε_{it})	3.49	3.36
	(0.000)	(0.000)
ρ	0.33	0.33
	(0.000)	(0.001)

Note: Model also includes year dummies.

Source: Authors' calculations using earnings records and HRS data; see text.

---

[21] The results are from a right-censored Tobit of the natural log of earnings with panel-level random-effects. The panel-level variance is Variance(v_i) and the overall variance is Variance(ε_{it}). Rho (ρ) is the ratio of the panel-level variance and the sum of both the panel-level and overall variances. When ρ is zero, the panel-level variance component is unimportant and the results should be identical to the results from a pooled estimator. Our results suggest that the panel-level variance is meaningful to the estimate of earnings. Additional detail on our estimates is available on request.