Wealth Inequality and Retirement Preparedness: A Cross-Cohort Perspective

John Sabelhaus
Alice Henriques Volz

October 2021

Abstract

High and rising U.S. wealth inequality underscores the need to revisit a perennial concern in policy circles: retirement preparedness. Our cross-cohort approach to studying retirement adequacy is based on *relative* wealth measures, meaning how the wealth distribution of one cohort compares to the cohorts ahead of them at the same age. We introduce *relative rank distributions* that show where individuals are in terms of the cohorts ahead of them at the same age, and *percentile point comparisons* that show how wealth levels at various percentiles vary across cohorts by age. We find that early Boomer's wealth is generally on par with or above 1930s cohort wealth at age 60. There is, however, evidence of relative wealth declines in the bottom of the wealth distribution for mid-late Boomers and Gen-Xers relative to earlier cohorts at younger ages, which is consistent with rising wealth inequality across and within generations. Social Security is an important offset to relative wealth declines at the bottom of the wealth distribution, but those benefits are not expected to be fully payable for the youngest cohorts.

JEL Codes: D15, G11, J26

John Sabelhaus is a Non-Resident Senior Fellow at the Brookings Institution and Adjunct Associate Research Professor at the University of Michigan. Email: jsabelhaus@gmail.com. Web: www.johnsabelhaus.com. Alice Henriques Volz is a Principal Economist at the Federal Reserve Board. Email: alice.h.volz@frb.gov. This paper was prepared for the Stanford Institute for Economic and Policy Research (SIEPR) 2019 Working Longer and Retirement Conference, October 14 -15, 2021. The analysis and conclusions set forth are those of the authors and do not necessarily reflect the views of the Brookings Institution or its funders or indicate concurrence by other members of the research staff or the Board of Governors of the Federal Reserve System.

1. Introduction

Much of the empirical research on the distribution of household wealth is focused on wealth inequality. A common question in many of those studies involves "top shares" of household wealth (Saez and Zucman, 2016; Bricker, Henriques, Krimmel, and Sabelhaus, 2016; Smith, Zidar, and Zwick, 2019). Although there is still substantial debate about exactly how much wealth is owned by the top 1 percent, top 0.1 percent, or the top 0.01 percent of the population, there is general agreement that household wealth is highly skewed and has become more unequal over time. Our recent paper shows that those conclusions about trends in top wealth shares hold even after broadening the concept of the wealth to include the present value of private and public retirement benefits (Sabelhaus and Henriques Volz, 2021).

Another important reason to conduct empirical research with household wealth data is to study lifecycle wealth accumulation, with a focus on the state of retirement preparedness or "adequacy." Some retirement adequacy studies compare accumulated wealth to the predictions of a lifecycle consumption smoothing model. Other studies involve more straight forward questions and calculations, such as whether the annuitized value of accumulated wealth is sufficient to maintain pre-retirement income or consumption levels. The literature on retirement adequacy is in many ways much less settled than the top wealth shares literature, with conclusions ranging from "more than half" to "fewer than 20 percent" of U.S. families are not adequately prepared for retirement.

One common element in both wealth distribution literatures is the need for a comprehensive measure of household wealth. The starting point for wealth is marketable assets less liabilities, which is the concept captured in the Survey of Consumer Finances (SCF) and estimated by the Federal Reserve Board at $96.1 trillion in the 2019 survey (Bhutta, et. al., 2020). However, household wealth also includes the present value of Defined Benefit (DB) pension claims, which adds another $19.1 trillion to household wealth in 2019. DB wealth is not measured directly at the household level, but the aggregates can be allocated across the individuals in the SCF who have legal claims to those benefits (Sabelhaus and Henriques Volz, 2019 and 2021). Finally, although the aggregate net present value of Social Security benefits less taxes (or Social Security Wealth, SSW) is just the relatively small balance in the Social Security trust fund, the net present value for working age and retired individuals is substantial, adding another $23.9 trillion to household wealth in 2019 (Sabelhaus and Henriques Volz, 2021).

Including DB pensions and SSW increases total household wealth by nearly 50 percent, but the distributional effect is disproportionate. DB wealth and SSW are more evenly distributed than the narrow marketable assets less liabilities concept captured directly in the SCF. Indeed, adding DB wealth and SSW to the SCF wealth concept lowers top 1 percent and top 10 percent wealth shares substantially in the 1995 to 2019 survey waves, depending on the specific sorting mechanism used to separate households or individuals into wealth groups (Sabelhaus and Henriques Volz, 2021). Including DB wealth and SSW does not change the trends in top shares, however, because the increasing wealth concentration within the narrow SCF wealth measure has not been offset by decreasing concentration within (or faster growing) DB wealth and SSW.

The comprehensive wealth concept that includes DB wealth and SSW can be measured over the lifecycle and time using the SCF in a pseudo-panel framework, and simple cross-cohort average wealth based on those data motivates our approach to studying retirement adequacy. As noted, the comprehensive wealth measure does not reverse the trend in top wealth shares over time, but it does reveal stark differences in average wealth across age groups over time (Sabelhaus and Henriques Volz, 2021). A disproportionate share of the growth in average comprehensive wealth between 1995 and 2016 occurred at older ages. Reducing those average wealth measures for expected Social Security funding shortfalls makes the age differentials larger, especially for the bottom half of the wealth distribution. In fact, average comprehensive wealth is estimated to be lower in 2016 than it was in 1995 for younger individuals.

In this paper we use the comprehensive wealth measures we developed for analyzing levels and trends in top wealth shares to study retirement adequacy. As noted, however, the issue of how to gauge the adequacy of retirement wealth is far from settled. One approach is to compare wealth to some target based on pre-retirement living standards, as measured by income or consumption. The simple financial advisor's rule that you should "replace 70 percent of your pre-retirement income" is often put forth as a straw man, but researchers then go on to develop more complicated calculations build on that same principle. An alternative approach is to compare individual wealth holdings against the predictions of a calibrated life cycle model. However, both approaches are sensitive to assumptions. What exactly does it mean to "maintain" pre-retirement living standards? What is the appropriate utility function for a calibrated life cycle model? What are the appropriate time preference and other parameters in such a model? How well does either approach capture environmental details like social insurance and income taxes?

We take an alternative approach to gauging retirement wealth adequacy that avoids these potential pitfalls. Our cross-cohort approach to studying retirement adequacy is based on *relative* wealth measures: how the wealth distribution of one cohort compares to the cohorts ahead of them at the same age. The relative wealth measures do not tell us anything directly about the fraction of a given population with or without adequate retirement wealth. However, given a reference point—say what fraction of current retirees are suffering hardship in retirement—we can draw conclusions about future retirees by looking at their *current* wealth distributions *relative* to the wealth distributions of current retirees when they were observed at younger ages.

Simply charting the wealth distributions for two or more birth cohorts at a given age does not tell us what we want to know about retirement adequacy, however. We introduce two different ways to look at relative wealth distributions. The first approach is *relative rank distributions* that answer the question, "Where would all of the individuals of a given cohort be if their wealth is mapped into the distribution of an earlier cohort?" The second approach is *percentile point comparisons,* that answer the question, "What are the wealth holdings for individuals at a given percentile of the wealth distribution across cohorts and ages?"

The relative rank charts tell us about people—how many individuals, in terms of comprehensive wealth, are ahead of or behind their counterparts in earlier cohorts at the same age. The key messages about inequality and retirement preparedness come from seeing the entire cohort arrayed along the (relative) wealth distribution. Complementing that perspective, the percentile points comparison charts tell us about dollars—cross-cohort differences in wealth at a given age for a fixed percentile of the distribution. The focus on one wealth percentile at a time allows us to drill down into the relevance of the various wealth components across the wealth distribution.

We find that early Boomer's (1940s cohort) wealth is generally on par with or above 1930s cohort wealth at age 60. That is, when we project the relative wealth positions of the 1940s cohort on the 1930s cohort wealth chart, we see that most 1940s relative percentile points are at or above the 1930s (reference) cohort at age 60. When we look at the percentile point charts, we see the 1940s cohort above the 1930s cohort at the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentiles in terms of the comprehensive wealth measure, and Social Security contributes importantly to that, especially at the bottom of the wealth distribution. This conclusion holds if

we compare wealth after adjusting for prices (using CPI) or for nominal wage growth (using the Social Security Average Wage Index).

Our data set spans 24 years, and thus we can also compare the 1950s cohort to the 1930s cohort at age 60. The relative wealth holdings of the 1950s cohort (the mid-Boomers) at age 60 gives us the first indications of slippage in relative wealth at the bottom of the wealth distribution. The extent of the slippage depends on the wealth concept and the approach to making wealth comparable across cohorts. The narrow concept (excluding SSW) adjusted for AWI shows (for example) that individuals in the 1950s cohort are 5 to 10 rank points behind their 1930s counterparts in the bottom half of the wealth distribution. The more comprehensive wealth measure including SSW and CPI adjustment reverses or shrinks those gaps considerably.

Looking at younger ages, and again depending on exactly how we compare wealth across cohorts, there is more evidence of relative deterioration in the bottom and middle of the wealth distribution for late Boomers and Gen-Xers. The gaps in the relative rank distributions are both larger and reach further up into the wealth distribution, and the percentile point comparisons show wealth levels below reference cohorts using the narrow wealth measure. For example, conventionally measured wealth at the 10^{th} and 25^{th} percentiles of the late Boomer and Gen-X distributions is well below the wealth of earlier cohorts observed at the same ages.

Social Security is an important offset to relative wealth declines at the bottom of the wealth distribution. Indeed, adding SSW and recomputing wealth distributions based on comprehensive wealth reverses many of the points about relative deterioration at low wealth levels across cohorts and ages. In that sense, Social Security has become *relatively* more important in terms of total wealth for the bottom half of the wealth distribution in younger cohorts, the result of the higher expected benefits based on lower (relative) earnings and higher (relative) life expectancy. However, that fact that those Social Security benefits are not expected to be fully payable for the youngest cohorts under current law overrides the relative improvements in retirement adequacy.

The paper proceeds as follows. In the next section we discuss how we construct our comprehensive measure of wealth. In the third section we review the literature on measuring retirement adequacy and discuss how the required (and debatable) assumptions motivate our relative wealth approach. The fourth section presents the relative rank analysis, and the fifth shows percentile point comparisons. Section six concludes.

2. Data and Methods

Our approach to studying retirement adequacy requires a comprehensive person-level wealth measure across birth cohorts and at various ages. Achieving this goal involves starting with high-quality household level balance sheet data, then adding household-level estimates of SSW and DB wealth.[1] The micro data used here is the Survey of Consumer Finances (SCF) for 1995 through 2019. The SCF is focused on household balance sheets, and the survey also has extensive information about incomes, demographics, and labor force experiences.[2] The SCF is a series of cross-section snapshots, so calculating lifecycle values for Social Security taxes and benefits requires estimating lifecycle earnings for individuals and their spouse/partners. Estimating present discounted values for both SSW and DB pensions also requires individual mortality rates, which we differentiate by age, sex, income, and birthyear.

Estimating the present discounted value of retirement benefits is complicated for several reasons. One conceptual issue is whether to use "expected" versus "termination" benefits (Sabelhaus and Henriques Volz, 2021). The concept of termination value for SSW—what any given individual would receive if the system shut down today—is not well suited for the analysis here, because the assumptions about who would receive benefits under a terminated system have little real-world relevance. Expected SSW is the present discounted value of benefits less taxes conditioned on the individual working and paying taxes through a given expected stop work age (as captured in the survey) and receiving benefits as soon as possible (age 62 or the first year after labor force exit, whichever is later). The labor force and earnings history along with expectations data in the SCF are used in conjunction with typical lifecycle earnings patterns derived from linked longitudinal survey and administrative data in the Health and Retirement Study (HRS). The detailed assumptions needed to create required inputs for computing expected SSW in the SCF are discussed in earlier papers (Sabelhaus, 2019; Sabelhaus and Volz, 2021).

The measure of DB wealth consistent with our comprehensive household balance sheet is the present value of future DB benefits, which is equivalent to the value of the financial assets held now that will be liquidated over time to pay the promised stream of DB benefits when those liabilities come due. The SCF collects details about DB pension benefits in three different survey

[1] For a more detailed explanation of the methods used to construct comprehensive household wealth, see Sabelhaus and Henriques Volz, 2021.
[2] See Bhutta et al (2020) for a discussion of the SCF and the most recent results, for survey year 2019.

modules. The three categories cover DB benefits already being received, DB benefits associated with a past job where the known benefit amount will be received at a specific future date, and DB benefits associated with a current job, where the ultimate benefit will depend on how much longer the worker is covered by the plan and their final salary.

For currently received benefits and reported expected benefits from past job pensions, the respondent is asked how much is currently being received, or how much will be received when the benefit begins. For our purposes, these streams of benefit payments are the input to a present discounted value calculation that also involves an assumed interest rate and (demographically differentiated) mortality (more on this below). This present discounted value calculation is the level of financial assets that the retirement plan sponsor must hold to pay those promised benefits and, thus, corresponds directly to the household wealth measure we are trying to capture.

Calculating DB pension wealth for workers covered by a plan on their current job is more complicated, because the benefit that will eventually be received is unknown as of the survey date. The measure of DB wealth that corresponds to our comprehensive total household wealth is the narrower financial liability of retirement plan sponsors: the DB wealth that the worker has accumulated to date. Plan sponsors are required to hold only the present value of benefits already earned by the worker, the termination value. The termination value represents the worker's legal claim to DB wealth because that level of assets is equivalent to the present value of benefits they will receive if their plan coverage ended today.

The core demographics in the SCF (age, sex, and income) are key inputs to estimating present values for retirement income streams through a differential mortality adjustment.[3] Constructing the present value of Social Security and DB pension incomes requires survival probabilities, which are computed for each SCF respondent and spouse/partner through age 99. The starting point for the survival adjustments used here to transform data from the cross section to the longitudinal structure is Social Security Administration cohort mortality by age and sex. The second step is a differential mortality adjustment by income percentile within age and sex groups based on Chetty et al (2016), based on income tax records linked to Social Security death records (for details, see the appendix to Sabelhaus and Henriques Volz, 2021).

[3] The demographic variables are available for both the SCF respondent and the spouse/partner, which is important, because the SCF micro files always assign the male in a couple (person closest to age 40 in a same sex couple) to be the respondent, and the other individual to be the spouse/partner, without regard to which is the primary earner.

3. How Do We Know If Retirement Wealth Is "Adequate"?

Financial advisors have long advocated a "replacement rate" approach to retirement planning. A standard rule of thumb is that individuals should be able to replace something like 70 percent of their pre-retirement income through their retirement years. The 70 percent rule is an average, and good financial advisors are quick to note that differences in individual circumstances will move the target up or down. Heterogeneity in circumstances is also a recognized key in the academic literature on retirement adequacy.

The question of retirement adequacy is far from resolved. One prominent academic "index" of retirement adequacy uses the SCF to estimate what fraction of the population is expected to have enough wealth at age 65 to maintain their pre-retirement standard of living, as measured by consumption (Munnell, Chen, and Siliciano, 2021). The authors find that about half of current U.S. households will fall more than 10 percent short of reaching their target, up from about a third of households in the 1980s. The model underlying the index is very complex, with income replacement targets that vary by factors such as housing tenure and effective income tax rates. However, key assumptions on inputs such as consumption levels during retirement are not consistent with observed age-spending patterns, and that biases absolute adequacy measures towards shortfalls (Hurd and Rohwedder, 2012).

Comparing actual wealth from a survey such as the SCF to predicted "optimal" wealth from a calibrated lifecycle model is an alternative way to measure retirement adequacy (Engen, Gale, and Uccello, 1999). Rather than ask what level of wealth is needed to meet a given target replacement rate, this approach defines "adequate" as enough wealth to smooth the predicted marginal utility of consumption over the lifecycle. Relatively simple lifecycle models can provide a wide range of predictions about optimal lifecycle wealth, depending on assumed inputs such as time discount factors, earnings paths, mortality, and even the arguments in the utility function. A lack of heterogeneity in structural model inputs leads to predictions that average across observations, and a larger fraction of the population will far short of those average targets.

Versions of the structural lifecycle model "optimal" wealth approach that capture more heterogeneity across the population generally seem to suggest that fewer U.S. families have inadequate retirement saving. One such study that uses Health and Retirement Study (HRS) longitudinal data allows the authors to make use of very detailed earnings histories and demographic characteristics and thus fine tune the lifecycle model predictions to specific types

of households (Love, Smith, and McNair, 2008). They find, for example, that only 18 percent of households would fall short of maintaining income at more than 150 percent of poverty over their remaining lifetimes.

Another, even more finely tuned comparison of actual and optimal wealth solves a different lifecycle optimization problem for every household in the HRS based on their unique characteristics (Scholz, Seshadri, and Khitatrakun, 2006). They find that fewer than 20 percent of households are saving below their optimal target, and the shortfalls of those who are under saving are generally small. The predicted wealth from any structural lifecycle model still depends on exactly which circumstances (or "state" variables) are included in the model, and computational constraints are limit which characteristics are included. In addition, structural models are inherently static in nature, and do not capture how circumstances might be evolving for future populations.

Dynamic microsimulation is another approach to capturing heterogeneity and does not suffer from the computational limits and static nature of structural lifecycle models. A dynamic microsimulation uses stochastic transition equations to age a population forward through time, simulating the wide range of outcomes that will be experienced. One such study projected future earnings and retirement incomes for the baby boom generation and found that typical outcomes should continue to improve for the baby boom relative to their parents, but changing demographic and earnings patterns are leaving more individuals in an economically vulnerable position (Butrica, Iams, and Smith, 2007). The dynamic microsimulation draws attention to retirement wealth adequacy for specific groups whose circumstances are changing over time, especially divorced women, never married men, Hispanics, high school dropouts, those with weak labor force attachment, and those with the lowest lifetime earnings.

Another recent paper looking at future retirement outcomes captures heterogeneity and many of the benefits of dynamic microsimulation in a purely empirical way (Brown, Dynan, and Figinski, 2020). The approach avoids lots of assumptions, and in many ways motivates our approach to studying retirement adequacy. The authors use early waves of the HRS, for whom we see economic and demographic characteristics just before and through retirement, to assess which characteristics of pre-retirees are useful for predicting economic hardship during retirement. They then use those correlations to predict how today's cohorts approaching retirement will fare in their retirement years. They find that the younger cohort is indeed more

likely, on average, to experience economic hardship, particularly men. This comparison of cohorts at younger ages as a predictor of future outcomes is close to what we do here. That is, what can we say about the likely outcomes for a pre-retiree group by looking at a currently retired group when the current retirees were younger?

In sum, the wide range of opinion in the literature about how to map observed wealth into measures of retirement adequacy leads to a wide range of conclusions about what fraction of U.S. households are prepared for retirement. One can look at the same individuals with the same wealth and arrive at very different conclusions about their retirement wealth adequacy. Indeed, the uncertainty about assumptions needed to create such "absolute" measures of retirement wealth adequacy directly motivates the alternative approach we use in this paper. Our starting point is current retirees, and the investigation focuses on how the wealth of younger cohorts compare to those current retirees when they were at the same age in some earlier year.

Thus, rather than ask whether a given family or individual has sufficient wealth for a secure retirement, we create "relative" adequacy measures by looking across cohorts at various ages. We introduce two ways to characterize relative wealth distributions, which answer the following questions. First, "Where would the individuals of a given cohort be if their wealth is mapped into the distribution of an earlier cohort observed at the same age?" We call this the *relative rank* distribution. The interpretations are all about *counts* of individuals. For example, if the relative rank distribution for a comparison cohort lies always at or above the reference cohort (which we show as a 45-degree line) we can say the comparison cohort is just as well or better off than the reference cohort at every point in the wealth distribution.

The second measure answers the question, "What is the wealth holding for the individual at a given percentile of the wealth distribution across cohorts and ages?" We call these charts *percentile point comparisons*. The percentile points comparisons do not show us the entire distribution like a relative rank chart. Rather, the percentile points comparison charts focus on a single slice of the wealth distribution across cohort and age charts. Another way to think of this: The relative rank charts show us counts of people at points in the wealth distribution and the percentile points comparisons are about dollar gaps for the people at those points. Though we can only focus on one wealth slice at a time, the distinct advantage to the percentile points is being able to stack and disentangle the contribution of different wealth sources at various percentiles of the wealth distribution.

4. Relative Rank Distributions

The first approach to using the comprehensive wealth data set to look at relative wealth distributions is what we refer to as *relative rank distributions*. These charts answer the question, "Where would the individuals of a given cohort be if their wealth is mapped into the distribution of an earlier cohort?" The distributional comparisons are specific to a given age group, a given wealth concept, and a given method for adjusting the wealth concept for comparability across the cohorts observed at the given age. Thus, there are many possible permutations for these sorts of relative rank calculations.

We limit the number of relative rank distribution charts by working with two wealth concepts across three age groups. SCF net worth plus DB benefits is the base wealth measure, while comprehensive wealth—the base measure plus SSW—is the alternative. The three age groups are 40, 50, and 60, with reference cohorts 1930s, 1940s, and 1950s. Given the time span covered by our SCF data set, the age 60 relative rank analysis (for example) uses 1930s as the reference cohort and 1940s and 1950s are the two comparison groups.[4] The age 50 and age 40 relative ranks shift reference and comparison cohorts forward one and two decades respectively.

The overarching goal of the relative rank (and percentile point) analysis is to avoid making absolute statements about what "adequate" retirement wealth means, but the approach does require an assumption about what it means for wealth levels to be "comparable" between cohorts. We are looking at reference versus comparison cohort wealth distributions 10 and 20 years apart in time, so simply comparing nominal wealth would be misleading. The question is thus whether the adjustment should consider only inflation, or some scalar that also involves changes in lifetime resources. For example, a cohort that had 10 percent higher real wages at every age might be expected to have 10 percent higher wealth at every age to maintain the same relationship between wealth and lifetime earnings. Our solution is to look at both inflation and nominal wage adjusted wealth distributions in the relative rank charts. Both the inflation and nominal wage adjustments are based on the data series used to calculate and inflation-adjust Social Security. The inflation adjustment uses the CPI, and the nominal wage adjustment uses the Average Wage Index (AWI).

[4] Technically, because the SCF data spans 1995 to 2019, we use the last seven years of each decade to represent a birth cohort, so the "1930s" means 1933-39 and "1970s" is 1973-79. Birth years within a given cohort are systematically observed with different frequencies across survey waves at various ages, so we reweight by birth year within each cohort and age group to keep the cohort representation constant across the relative comparisons.

Figure 1. Relative Rank Distributions at Age 60, Net Worth Plus DB Wealth

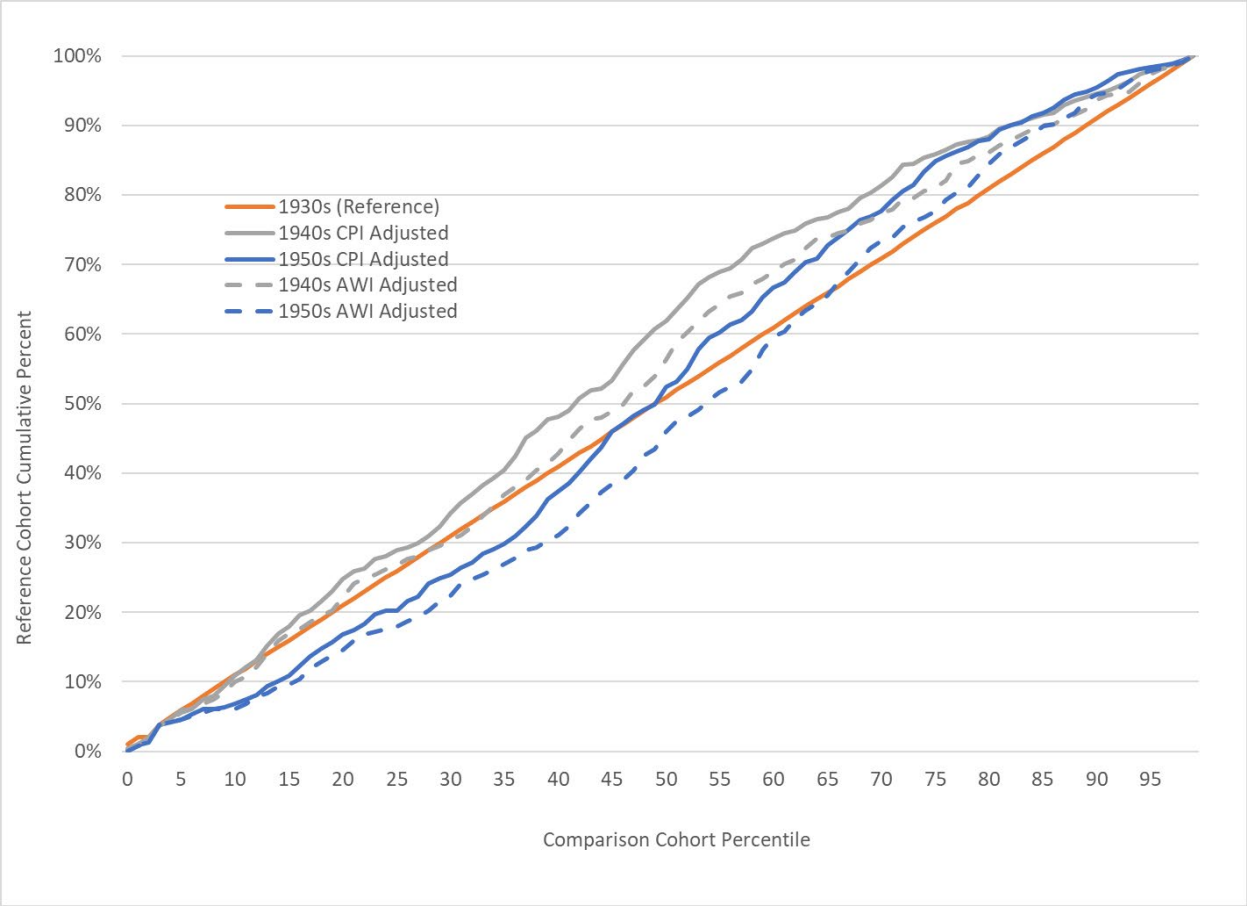

The relative rank distributions plot the comparison cohort percentiles on the x-axis, and the reference cohort cumulative percentiles on the y-axis. The graphs all include a 45-degree line that plots the reference cohort cumulative wealth against their own wealth distribution. Thus, the 45-degree line (trivially) shows that one percent of the reference cohort population is in each percentile of the reference cohort wealth distribution. However, that 45-degree line is a useful benchmark for the comparison cohorts, because we can say that if the comparison cohort relative rank distribution falls along the 45-degree line, the comparison cohort has the same wealth distribution as the reference cohort. If the comparison cohort relative rank distribution is above the 45-degree line, comparison group wealth is above reference cohort wealth at that percentile.

With those chart-reading principles in mind, we can look at our first relative rank distribution, using the base wealth concept (SCF net worth plus DB wealth) at age 60 (Figure 1). The reference cohort for age 60 is the 1930s. The grey lines show the relative rank distributions for the 1940s (loosely, early-Boomers) cohort, and the blue lines show the relative rank

distributions for the 1950s (loosely, mid-Boomers). The solid lines are CPI-adjusted wealth, and the dashed lines are AWI adjusted wealth.

Both the CPI-adjusted and AWI adjusted relative rank distributions for the 1940s cohort are on or above the 45-degree line. That means everyone in the 1940s cohort had the same or more wealth than their counterpart at the same point in the wealth distribution in the 1930s cohort at the same percentile. Both the solid (CPI) and dotted (AWI) lines for the 1940s group are at or above the 45-degree line. Thus, the statement about relative ranks holds even if we adjust the wealth distribution percentile points such that the 1940s group is required to have the same or greater wealth even after adjusting for nominal wage growth.

The relative rank distributions only show people, not dollars. There is no statement one can make about differences in the dollar wealth distributions—we discuss that in the next section. However, the relative rank charts do allow us to make statements about where individuals' wealth levels lie relative to the reference cohort. The horizontal distance between the relative rank and 45-degree lines is that mapping. For example, focusing in on the 70th percentile along the 45-degree line, the CPI-adjusted relative rank for the 1940s cohort is around 55. That means an individual around the 55th percentile of the 1940s cohort has the same wealth as an individual at the 70th percentile of the reference cohort. They are about 15 rank points ahead. On an AWI-adjusted basis, the 60th percentile of the 1940s group corresponds to the 70th percentile of the 1930s distribution, meaning they are about 10 rank points ahead.

Although the relative rank distributions show an unambiguous improvement between the 1930s and 1940s wealth distribution at age 60, the story is different for the 1950s cohort relative to the 1930s. The 1950s group tracks the first few percentiles of the 1930s cohort, but then falls behind by a few rank points for much of the distribution. The specific relative rank gaps depend on whether one looks at CPI-adjusted or AWI-adjusted relative wealth, but somewhere between 40 and 60 percent of the 1950s cohort has less wealth than their counterparts in the 1930s group. The fact that the 1950s cross the 45-degree line at higher percentiles is consistent with what we know about rising wealth inequality over time. Again, however, there are no dollar amounts in the relative rank charts. All we can say is that an individual in the 1950s cohort at (say) the 25th percentile of their wealth distribution has the same wealth as someone in the 1930s at the 20th percentile, which (although no dollars are attached) seems like a modest gap at most.

Figure 2. Relative Rank Distributions at Age 60, Net Worth Plus DB and SSW

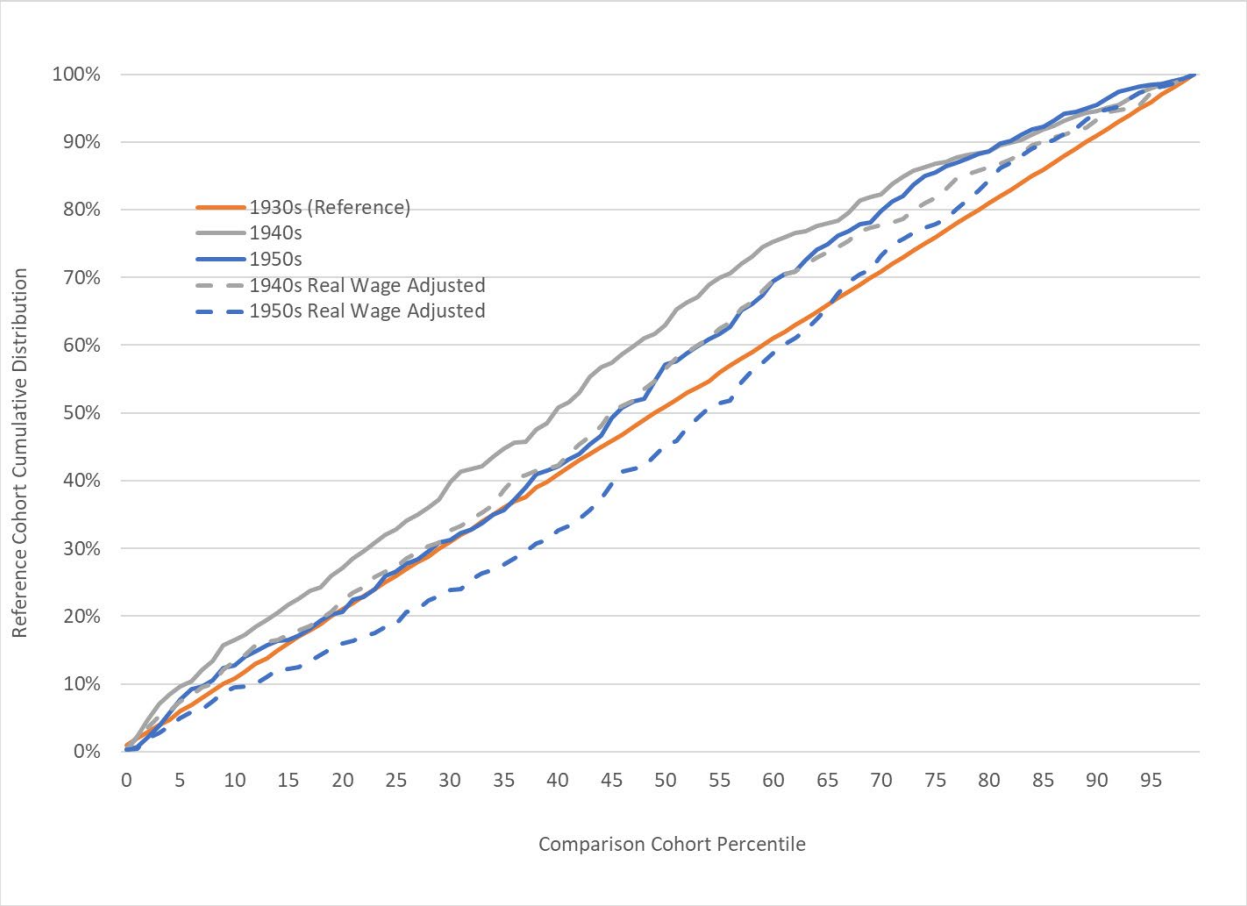

Adding expected SSW to the wealth concept pushes the 1940s cohort even further above the 1930s, and it improves the relative ranks for the 1950s cohort as well (Figure 2). One way to interpret this shift in relative ranks is that Social Security is *relatively* more important for the 1940s and 1950s comparison cohorts than it was for the 1930s reference cohort. The approach we use to estimate SSW in the micro data captures the fact that changing demographics and lifetime wages lead to differences in SSW for individuals at the same point in the wealth distribution across cohorts (see Sabelhaus and Volz, 2021, for details). If the typical individual in the comparison cohort at a given age and wealth percentile has longer life expectancy, they will have higher SSW. Changes in Social Security replacement rates, marriage patterns (through spousal and survivor benefits), and labor force participation also matter, but the relative ranks show that on net SSW is (again, relatively) more important for the younger cohorts.

Figure 3. Relative Rank Distributions at Age 50, Net Worth Plus DB Wealth

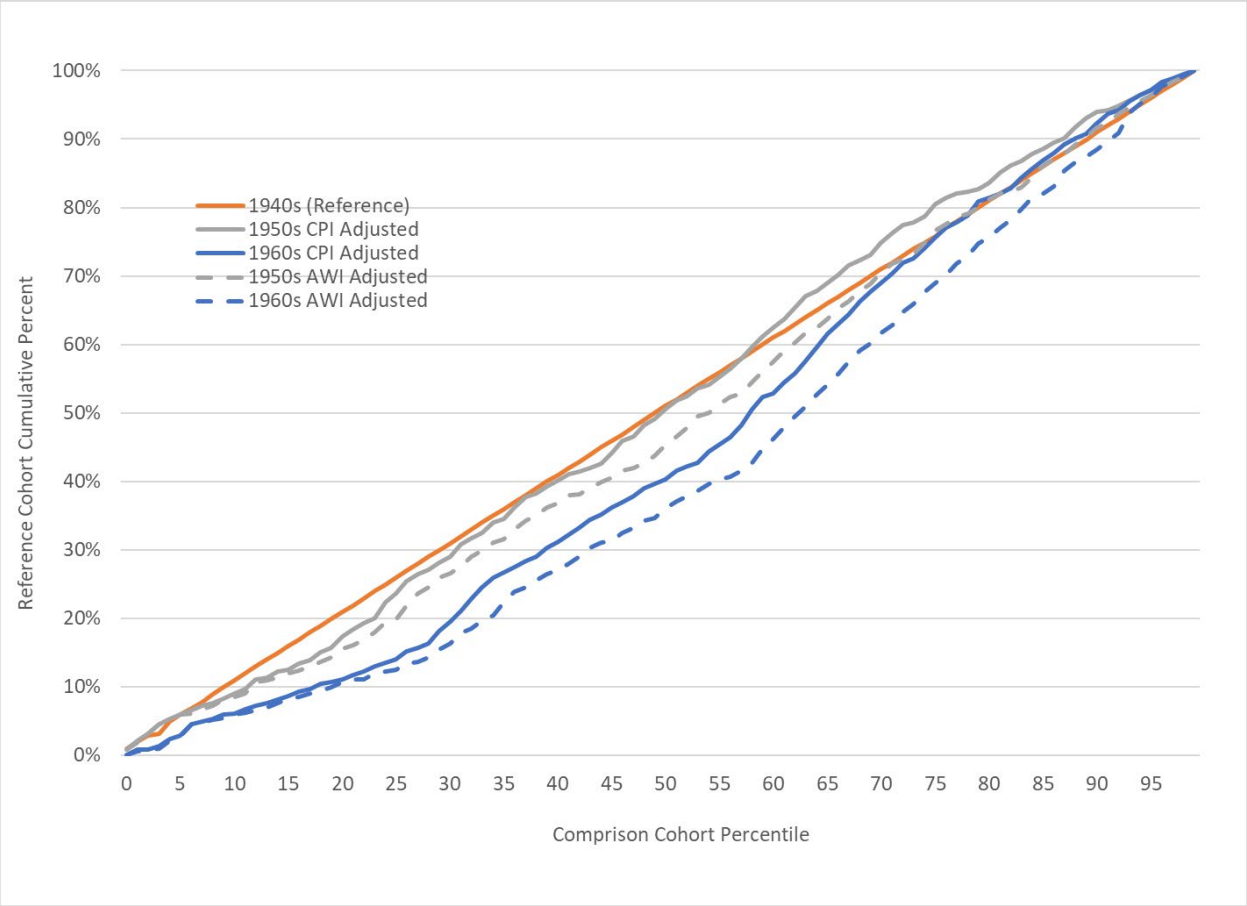

Shifting back to the base wealth concept without Social Security (SCF net worth plus DB wealth) the age 50 relative ranks show that the 1950s cohort is slightly trailing their early Baby Boomer counterparts at low to modest wealth levels, but then generally within a few rank points (Figure 3). The 1960s cohort trails the early Boomers by a much wider margin, on the order of 10 to 15 rank points over much of the distribution, depending on whether one uses the CPI or AWI comparability adjustment. The relative decline in wealth at low wealth levels between the 1940s and 1960s group is pronounced, with (for example) individuals around the 20th to 25th percentiles of the 1960s cohort having only as much wealth as an individual at the 10th percentile of the 1940s wealth distribution. In both the 1950s and 1960s comparison groups, the highest wealth individuals catch up to and surpass their counterparts in the 1940s (early Boomer) reference cohort.

Figure 4. Relative Rank Distributions at Age 50, Net Worth Plus DB and SSW

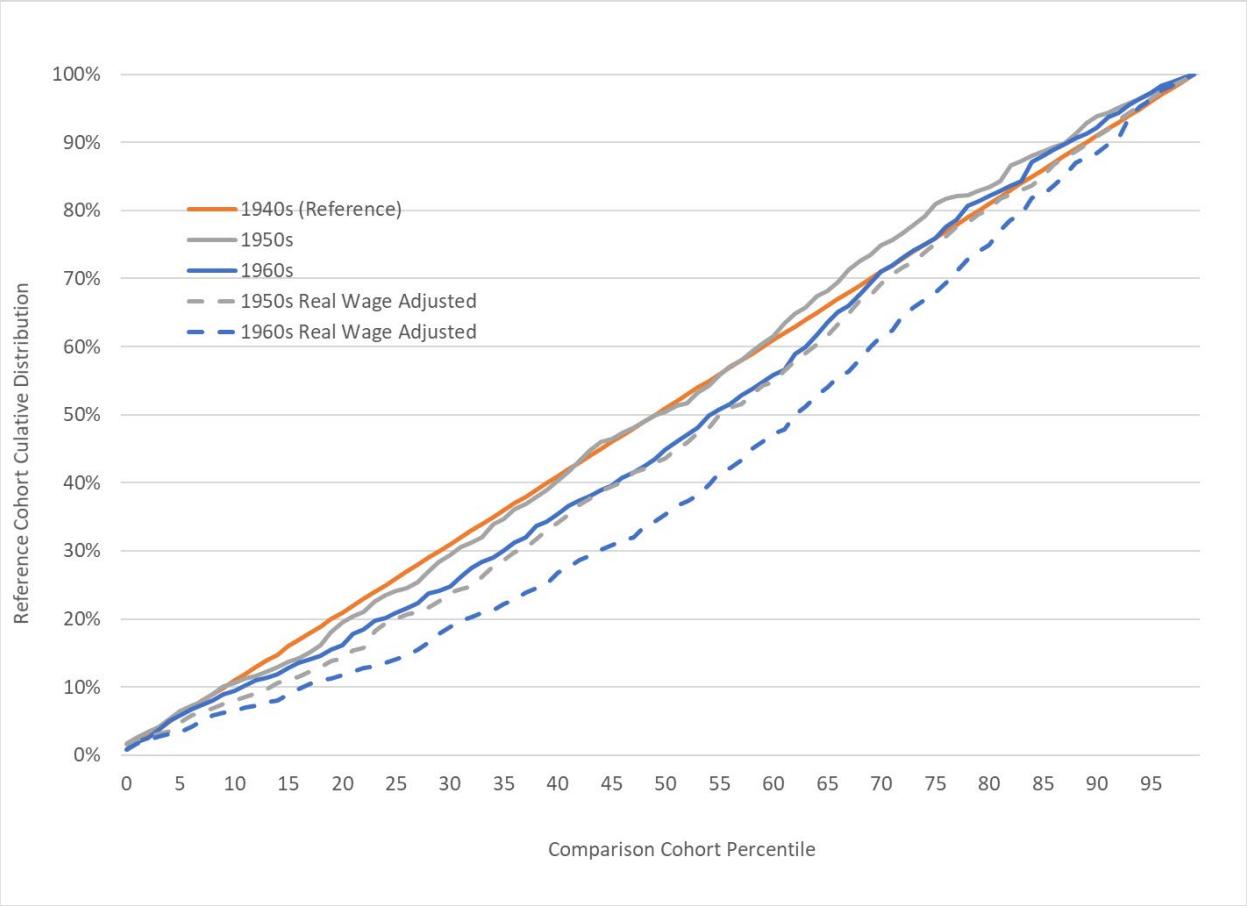

Adding SSW to the base wealth concept does not do much to the relative ranks at age 50 (Figure 4). As a reminder, the reference cohort generation wealth distribution also includes SSW, so the relative rank lines will only shift if Social Security is *relatively* more important for the comparison cohorts at a given wealth percentile. In this case, the relative ranks for the 1950s and 1960s comparison cohorts using the comprehensive wealth concept (SCF net worth plus DB and SSW) look very similar to the relative ranks using the narrower concept (SCF net worth plus DB) shown above (Figure 3). The normative takeaways are thus very similar, and to some extent dependent on how wealth is adjusted for comparability over time. The CPI-adjusted wealth percentiles suggest modest deterioration—a few rank points at most—between the 1940s and 1960s cohorts, though it is over most of the wealth distribution. The AWI adjustment sets a much higher bar for wealth comparability, and the relative rank gaps are as high as 10 or even 15 rank points.

Figure 5. Relative Rank Distributions at Age 40, Net Worth Plus DB Wealth

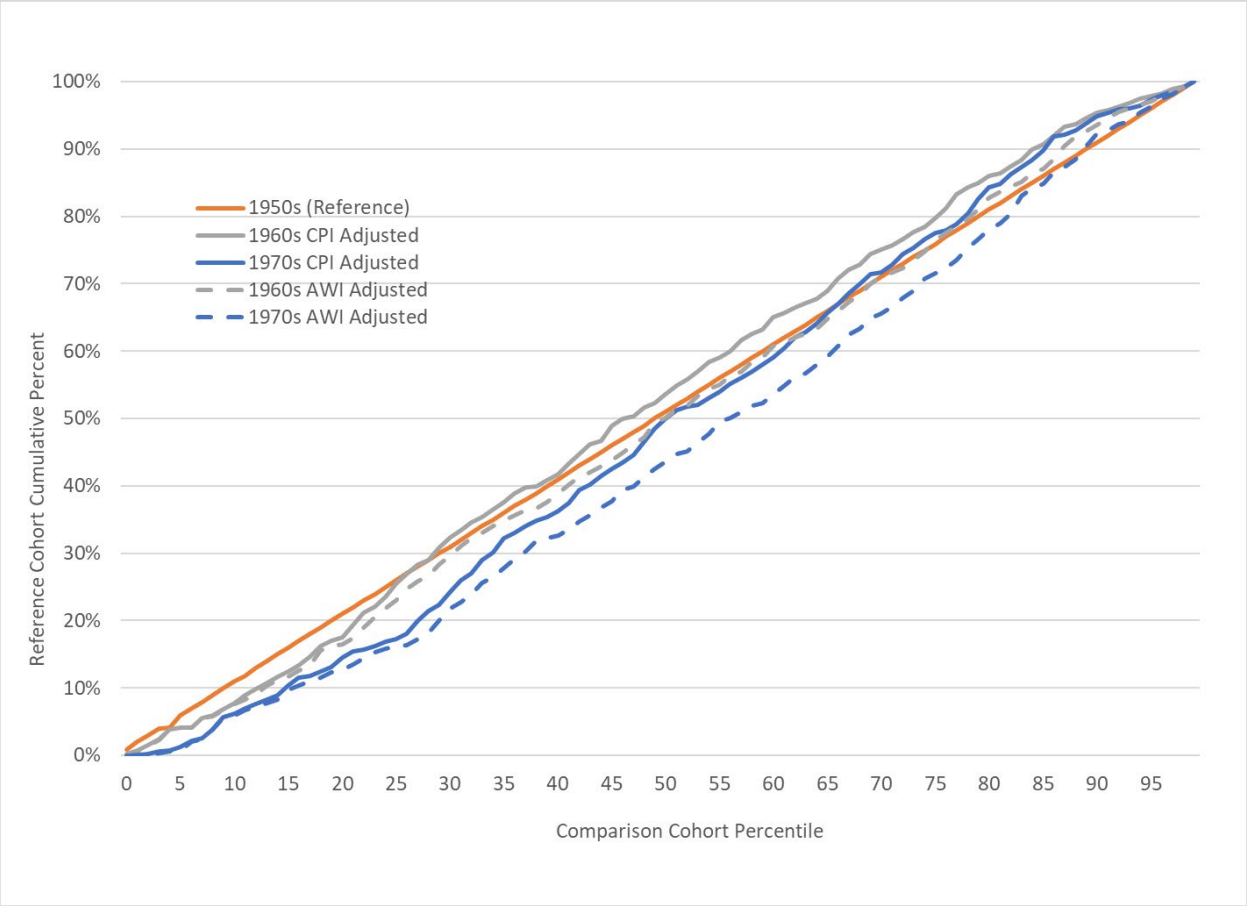

As noted in the introduction, our interest in relative wealth across cohorts at various ages is motivated in part by the observation that average wealth at younger ages is declining over time, which is another way of saying that a disproportionate share of the wealth gains have gone to older age groups (Sabelhaus and Volz, 2021). The approach here allows us to consider the average wealth declines at younger ages in two additional dimensions: where in the wealth dimension, and for which cohorts. The relative rank distributions for SCF net worth plus DB wealth at age 40 confirm relative wealth declines at low wealth levels for the 1960s cohort relative to the 1950s, and even bigger relative wealth declines for the 1970s cohort (Figure 5). It is worth noting here that much of the age 40 non-Social Security wealth distribution for all wealth levels is close to zero, a point that becomes clear when we look at low wealth percentiles in the percentile point comparison charts in the next section.

Figure 6. Relative Rank Distributions at Age 40, Net Worth Plus DB and SSW

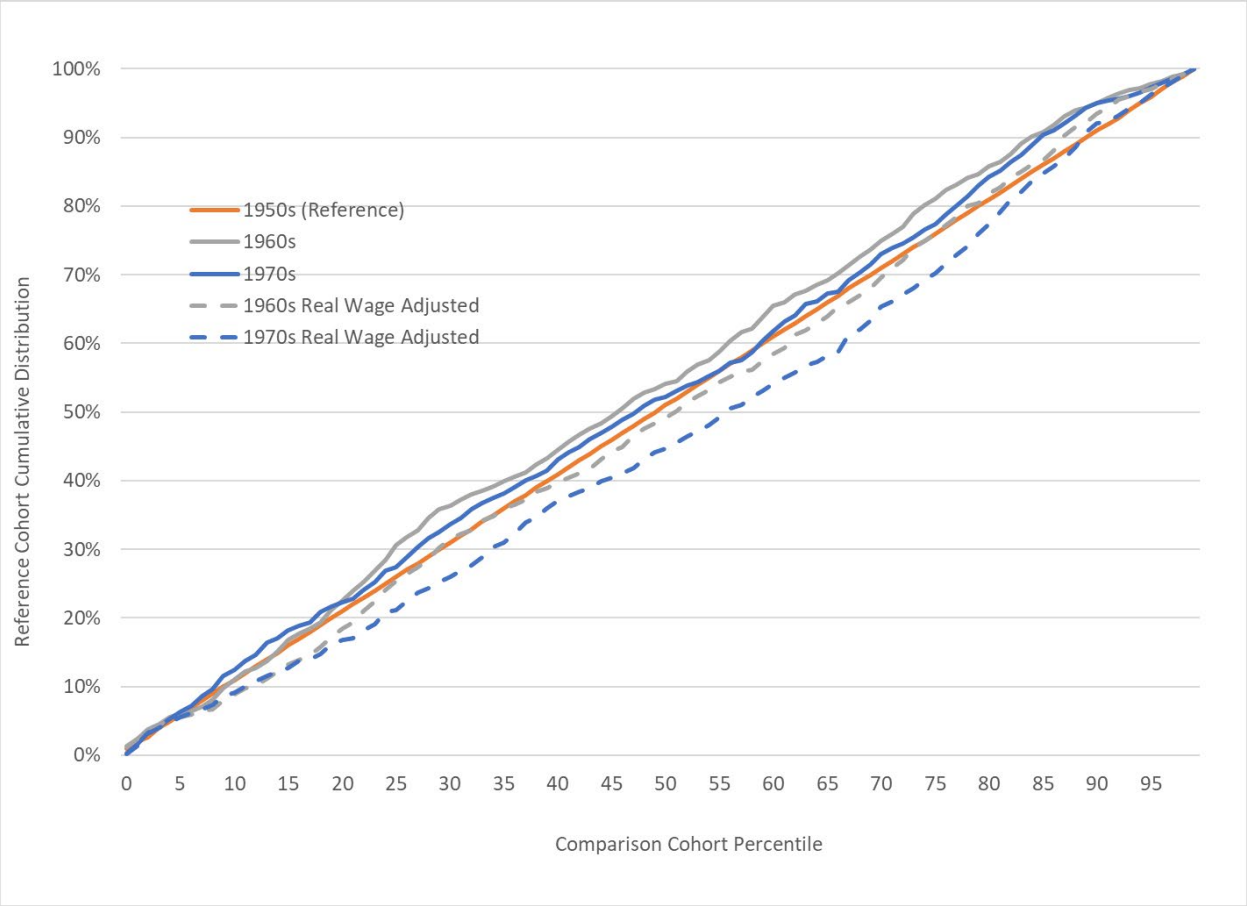

Social Security was shown to be relatively important for relative rank distributions at ages 50, and that holds at age 40 as well (Figure 6). Using a comprehensive wealth measure that includes all SSW—payable or not—the 1960s and 1970s cohorts are slightly above or slightly behind the 1950s group, depending on whether one uses the CPI or AWI wealth comparability adjustment. Indeed, tying together the relative rank distributions at various ages using the most comprehensive wealth measure (Figures 2 and 6) one can assert that (1) the 1950s cohort tracks or is slightly behind the 1930s cohort at age 60, (2) the 1970s cohort tracks or is slightly behind the 1950s cohort at age 40, and thus, if the relative ranks are transitive across time, (3) the 1970s cohort is somewhere between tracking and slightly behind the 1930s cohort.

Figure 7. Relative Rank Distributions at Age 40, Net Worth Plus DB and Payable SSW

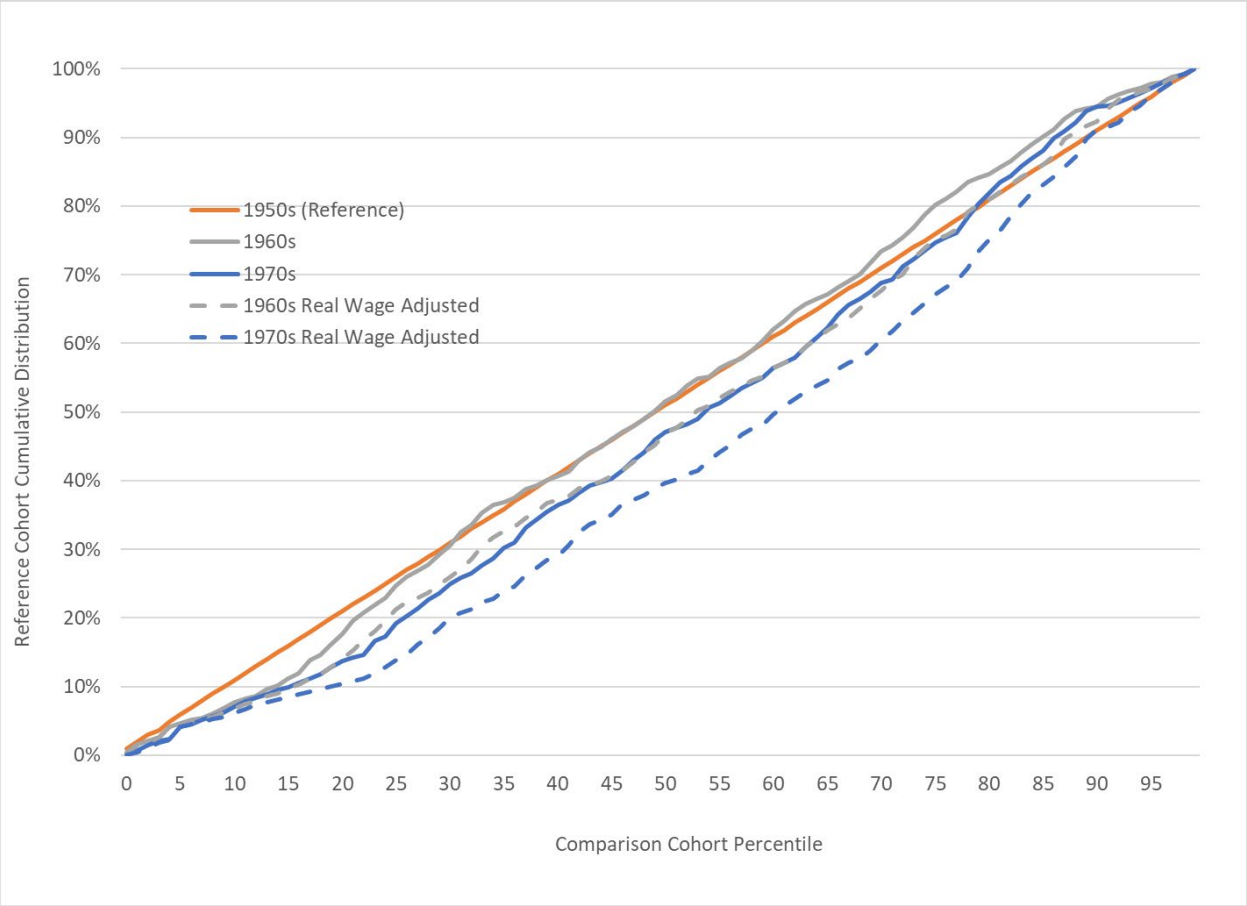

The optimistic relative rank conclusions at age 40 very much depend on whether the Social Security wealth offsetting relative declines in other types of wealth are in fact payable. Our assumption used to estimate payable benefits is simple: we adopt the Social Security Administration actuaries' projection that 80 percent of benefits are payable after 2035, and our payable SSW is the present value of scheduled benefits through 2034, and 80 percent of scheduled after that (see Sabelhaus and Volz, 2021). The timing is such that the youngest cohorts will face the largest lifetime cuts, and that comes through in the relative ranks. When we recompute relative ranks at age 40 using only payable SSW, there is evidence of relative slippage for the 1960s and 1970s cohorts across much of the wealth distribution (and remember, the 1950s reference cohort will also face some benefit cuts in the payable scenario, just proportionally less). But how much relative wealth decline does payable benefits suggest, and for whom? We turn to that in the next section.

5. Percentile Point Comparisons

Cross-cohort shifts in relative rank positions are useful because they describe one entire cohort group relative to another. The shortcoming is that we are unable to say anything about how large any shortfalls are, because the charts are ordinal in nature. Thus, our second set of charts—what we call *percentile point comparisons*, are designed to answer the question, "What is the wealth level for the individual at a given percentile in wealth distributions across cohorts and ages?" The percentile points comparison charts focus on a single slice of the wealth distribution across cohort and age groups. In addition to attaching dollars to the relative rank gaps, another advantage is being able to stack and disentangle the contribution of different wealth sources at various percentiles of the wealth distribution. Building on the relative rank charts in the previous section, we are particularly focused on three wealth components: SCF net worth plus DB, payable SSW, and scheduled but not payable SSW.

The percentile points comparison charts require a key statistical input. Although it is obvious what it means to be at a given percentile for a given wealth concept, working with multiple wealth components at a given percentile requires a decision about how to "stack" wealth measures to show the dollar contribution of each wealth component. One could in principle identify individuals at a given percentile of the base wealth concept (SCF net worth plus DB) and then add average payable SSW and average scheduled but not payable SSW. That approach introduces the complication of how to estimate the average incremental wealth for an individual at a given percentile, and in any case, does not generate the concept we care about, as reflected in the relative rank approach.

The alternative approach we take to building the percentile points charts begins with each wealth distribution separately, but we do not simply sum (say) the 10^{th} percentile of each wealth distribution to get an overall 10^{th} percentile. Rather, we start with total wealth (SCF net worth plus DB and all SSW) for a given percentile/cohort/age combination and solve for the 10^{th} percentile of that wealth measure. We then repeat the process for the other two measures: SCF net worth plus DB and payable SSW, and finally for just SCF net worth plus DB. The 10^{th} percentile of SCF net worth plus DB and payable SSW is then solved for as the difference between the first two percentile points, and the 10^{th} percentile of SCF net worth plus DB is the difference between the second and third percentile points.

Figure 8. 10th Percentile of Cross-Cohort Wealth Distributions

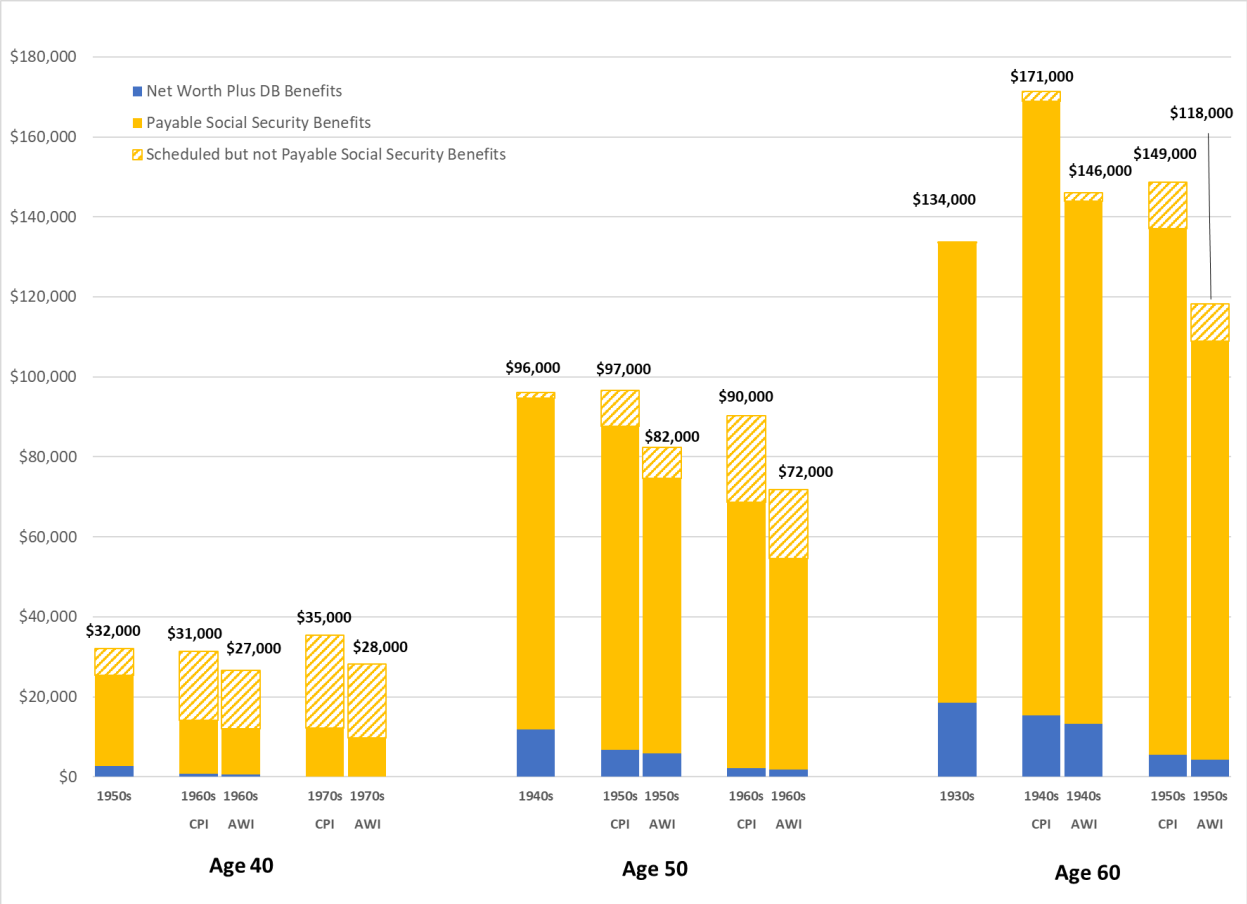

We begin the percentile comparisons near the bottom of the wealth distribution, at the 10th percentile (Figure 8). The differences in overall bar heights within a given age and cohort comparison are dollar-valued measures of the gaps in the relative rank charts at the 10th percentile using the most comprehensive wealth measure. The fact that overall bar heights are generally similar is consistent with the observation that relative ranks were generally close to the 45-degree line at the 10th percentile in the comprehensive wealth charts. One notable exception—the 1940s cohort at age 60 has more wealth at the 10th percentile than the 1930s cohort—is consistent with the observation that the 1940s relative rank line is above the 45-degree line at low wealth levels (Figure 2). The other standout feature of the 10th percentile point comparison is the role of SSW, first in overall dominance at every age, and second in terms of the importance of payable versus scheduled benefits at younger ages. The relative wealth of the 1970s cohort is dramatically lower than predecessor cohorts based on payable benefits alone.

Figure 9. 25th Percentile of Cross-Cohort Wealth Distributions

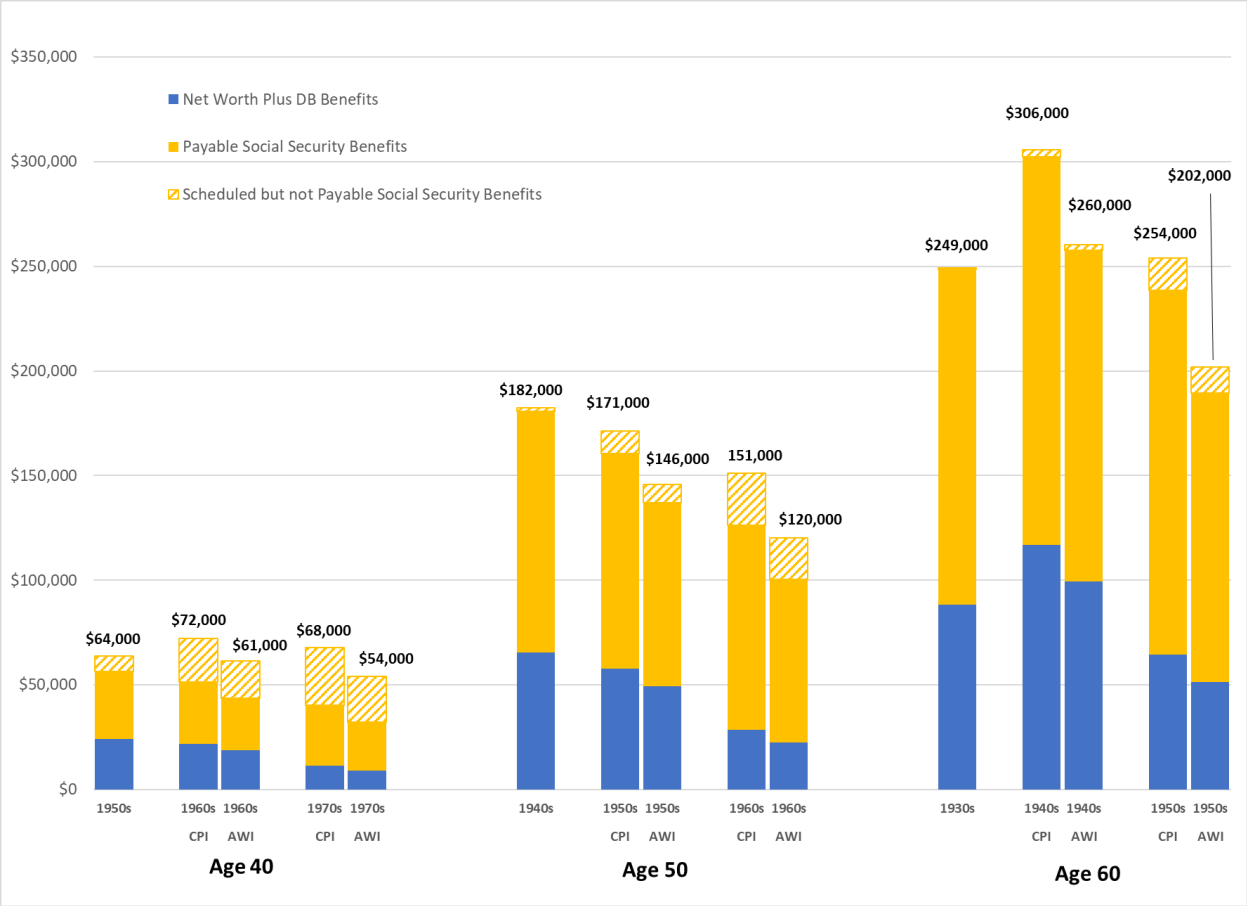

At the 25th percentile of the wealth distribution, base wealth (SCF net worth plus DB) begins to play a more important role, but Social Security still dominates (Figure 9). The relative deterioration in comparison cohort base wealth at the 25th percentile is again consistent with the gaps in the base concept relative rank charts (Figures 1, 3, and 5) at the 25th percentile. In most of those cases, the comparison cohort relative rank is below the 45-degree line at the 25th percentile. Again, one notable exception is the 1940s cohort at age 60—they have higher base wealth, which is the same as saying their relative rank is above the reference cohort at the 25th percentile. And as with the 10th percentile, the effect of Social Security is to reverse some of the relative wealth decline at the 25th percentile, but if one focuses on payable benefits alone, the prospects for low wealth individuals in (especially) the 1960s and 1970s cohorts are much bleaker.

Figure 10. 50th Percentile of Cross-Cohort Wealth Distributions

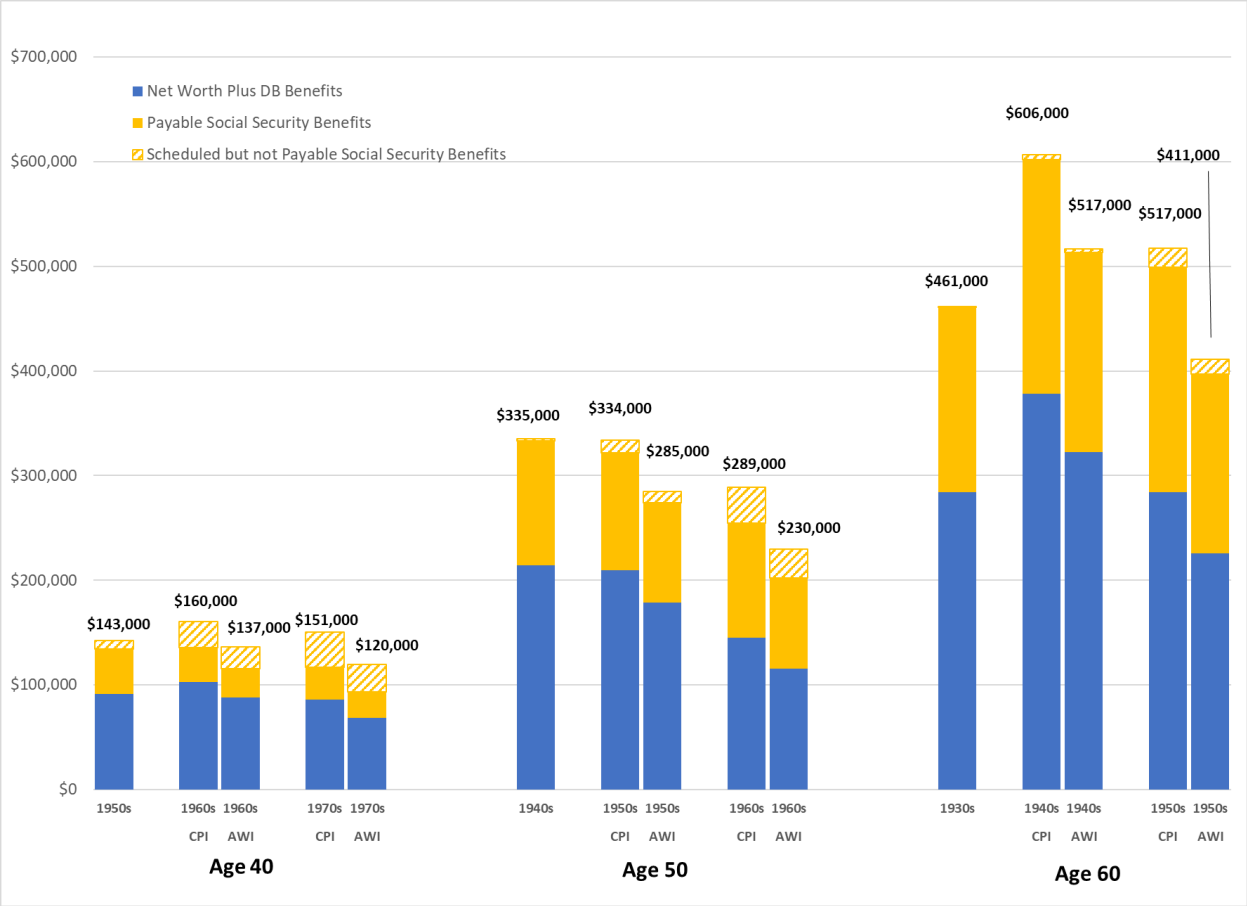

The contributions to total wealth from Social Security and non-Social Security sources are roughly balanced in the middle of the wealth distribution (Figure 10). The same observations about relative declines in base wealth (SCF net worth plus DB) we saw at the 25th percentile also show up at the 50th percentile, which is the same as saying the relative rank lines were generally below the 45-degree line over the bottom half of the wealth distribution. However, the relative changes are much more muted or go in the other direction using the comprehensive wealth measure, consistent with relative ranks above the 45-degree line at the 50th percentile (Figures 2, 4, and 6). As with the relative ranks, the specific inferences depend on whether one uses CPI or AWI adjusted wealth, but in general the 50th percentile total wealth measures show less relative deterioration across cohorts when compared to lower wealth levels.

Figure 11. 75th Percentile of Cross-Cohort Wealth Distributions

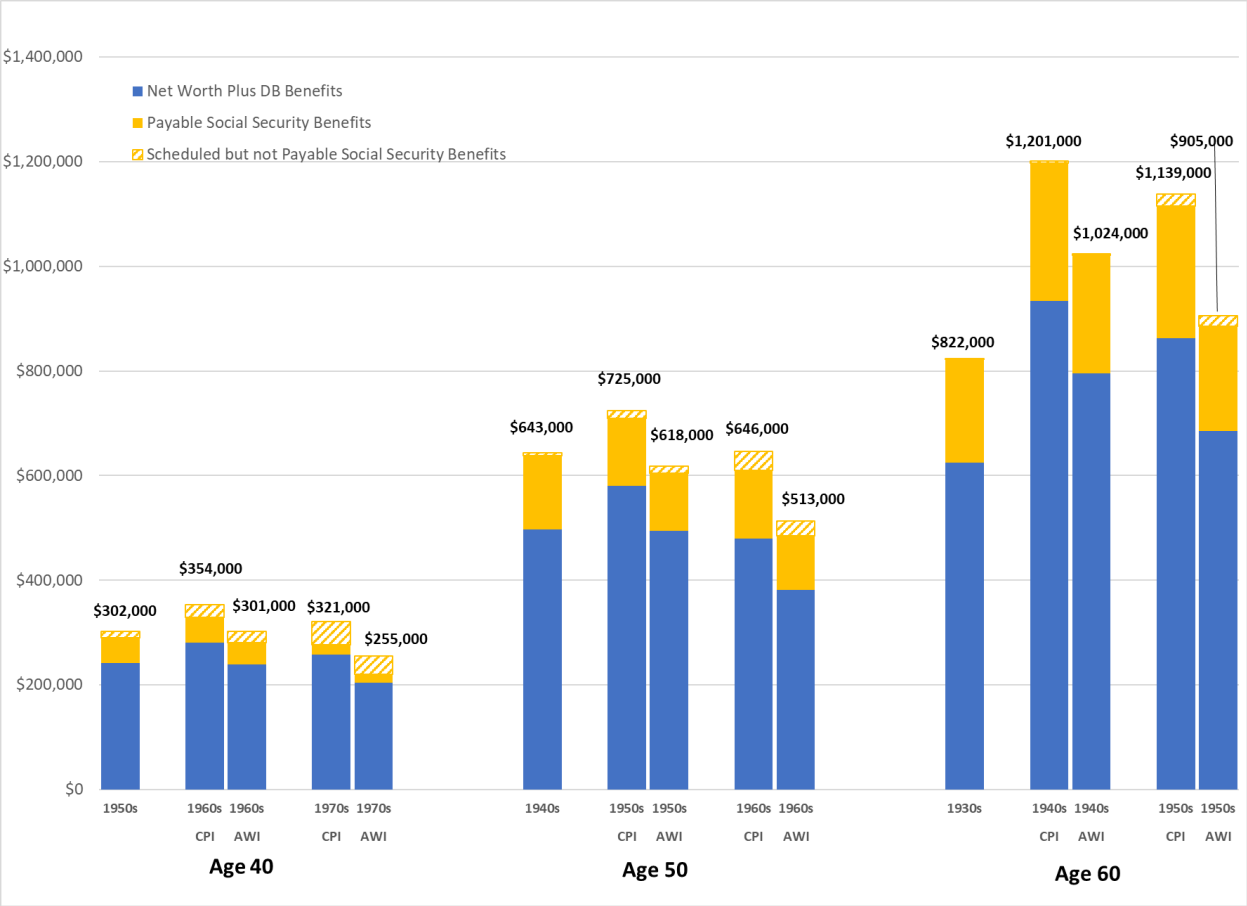

Retirement adequacy is about more than wealth levels, even in our relative wealth framework. Although it is possible that relative wealth would decline across cohorts at a given age in the top half of the wealth distribution, the data suggest that there is in fact less to worry about at the 75th percentile (Figure 11). In all cases based on CPI adjustments and total SSW (payable and not payable) the data show gains in relative wealth at the 75th percentile. Relative wealth declines are evident for younger cohorts at younger ages using the higher bar (AWI) for wealth comparability. Also, relative to the bottom half of the wealth distribution, Social Security is much less important for overall wealth in the top half of the wealth distribution. An important corollary to that statement is that Social Security solvency is—at least compared to the bottom half of the wealth distribution—a relatively less important determinant of the prospects for currently young higher wealth individuals.

Figure 12. 90th Percentile of Cross-Cohort Wealth Distributions

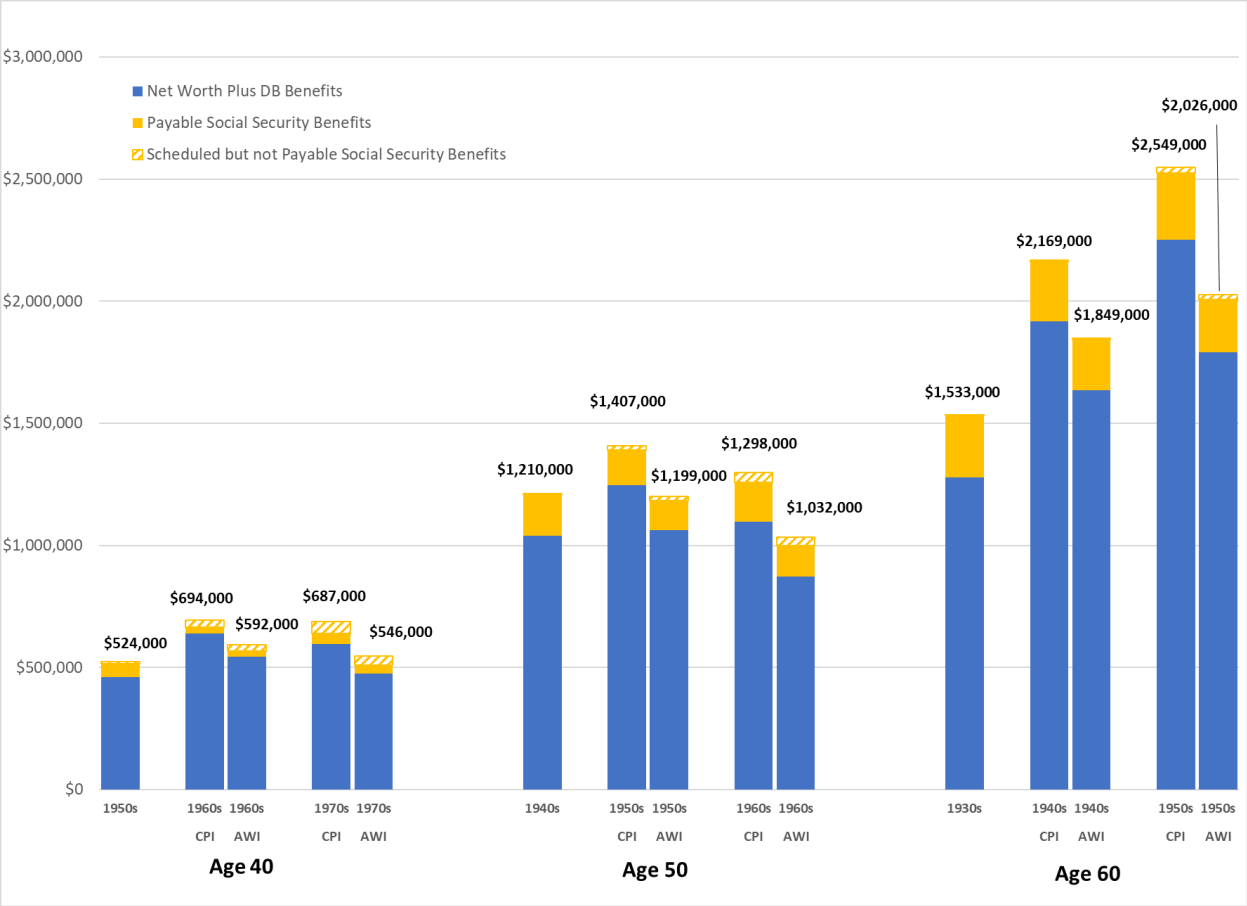

For completeness, we also show wealth components at the 90th percentiles of the wealth distribution across our age and cohort groups (Figure 12). The impression is the mirror image of the 10th percentile charts, in which Social Security dominates. At the 90th percentile Social Security is a relatively small component of total wealth, and the issue whether Social Security is payable or not has very modest implications for overall retirement wealth adequacy. It is worth noting that the 90th wealth percentile is in many ways like the 75th—younger cohorts are not unambiguously better off, and inferences about whether they are just as well or better off (versus a little worse off) depend on the method used to adjust wealth for comparability. The 90th percentile snapshot is also consistent with the idea that rising wealth inequality is more about what's happening within the top 10 percent, and not the 90th percentile relative to other groups per se.

6. Conclusions

Are today's 60-year-olds (born in 1960) well prepared for retirement? Are today's 50-year-olds (born in 1970) and today's 40-year-olds (born in 1980) on a trajectory to be well prepared when they retire? Against a backdrop of rising wealth inequality, the concern about *some* future retirees seems justified. But *how many* future retirees will face economic hardships in their retirement years? How *large* are the expected shortfalls in terms of wealth needed for a secure retirement? Two researchers with very different retirement wealth adequacy yardsticks can look at the same wealth distributions and come to very different conclusions about the number of people facing retirement shortfalls, and how large those shortfalls might be.

In this paper we consider retirement wealth adequacy using *relative* yardsticks. The reference points are today's retirees, the cohort born in the 1930s, the early Boomers born in the 1940s, and the mid-Boomers born in the 1950s, all of whom are in various stages of retirement. We compare the wealth distributions of younger cohorts at the same ages, and we find support for the idea that rising wealth inequality in net worth and DB pensions is indeed driving relative declines (and presumably retirement preparedness) in the bottom half of the wealth distribution.

The *relative* wealth approach means we cannot say anything about absolute retirement preparedness, for a given individual or a cohort in general. What we can do is make statements like the following. If we believe that (for example) 20 percent of the 1930s cohort experienced financial hardship in retirement, then some fraction below 20 percent of the 1940s cohort can be expected to experience financial hardship, based on their relative wealth distributions at age 60. Specifically, depending on the wealth concept and wealth comparability adjustment, we can say that somewhere between the 10^{th} and 15^{th} percentile of the 1940s cohort had as much wealth as the 20^{th} percentile of the 1930s cohort at age 60.

Also, and still depending on the specific wealth measure, the adjustment to make wealth comparable over time, and the assumption about future Social Security benefit cuts, the inferences go in the other direction for younger cohorts, especially at low wealth levels. Under a payable Social Security scenario based on the AWI wealth comparability adjustment, the 40^{th} percentile of the 1970s cohort has the same wealth as the 20^{th} percentile of the 1950s cohort at age 40, meaning we could be looking at double the fraction of retirees facing hardship if the 20^{th} percentile is our cutoff. At the bottom of the wealth distribution, relative retirement preparedness is very dependent on how policymakers address expected Social Security shortfalls.

7. References

Batty, Michael, Jesse Bricker, Joseph Briggs, Elizabeth Holmquist, Susan McIntosh, Kevin Moore, Eric Nielsen, Sarah Reber, Molly Shatto, Kamila Sommer, Tom Sweeney, and Alice Henriques Volz. 2019. "Introducing the Distributional Financial Accounts of the United States," Federal Reserve Board: FEDS Working Paper 2019-017. (March)

Bhutta, Neil, Jesse Bricker, Andrew C. Chang, Lisa J. Dettling, Sarena Goodman, Joanne W. Hsu, Kevin B. Moore, Sarah Reber, Alice Henriques Volz, and Richard A. Windle. 2020. "Changes in U.S. Family Finances from 2016 to 2019: Evidence from the Survey of Consumer Finances*," Federal Reserve Bulletin*, 106(5): 1-42. (September)

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus. 2016. "Measuring Income and Wealth at the Top Using Administrative and Survey Data," *Brookings Papers on Economic Activity, 1:2016*, p. 261-321.

Brown, Jason, Karen Dynan, and Theodore Figinski. 2020. "The Risk of Financial Hardship in Retirement: A Cohort Analysis," in Olivia S. Mitchell and Annamaria Lusardi eds., *Remaking Retirement: Debt in an Aging Economy*, Oxford University Press.

Butrica, Barbara, Howard M. Iams, and Karen E. Smith. 2007. "Understanding Baby Boomers' Retirement Prospects," in Brigitte Madrian, Olivia S. Mitchell, and Beth J. Soldo, eds., *Redefining Retirement: How Will Boomers Fare*, Oxford University Press.

Engen, Eric M., William G. Gale, and Cori E. Uccello. 1999. "The Adequacy of Household Saving," *Brookings Papers on Economic Activity, 1:1999*, p. 65-187.

Hurd, Michael J., and Susann Rohwedder. 2012. "Economic Preparation for Retirement," in David Wise, ed., *Investigations in the Economics of Aging,* University of Chicago Press.

Jacobs, Lindsay, Elizabeth Llanes, Kevin Moore, Jeffrey Thompson, and Alice Henriques Volz. 2019. "Wealth Concentration in the United States Using and Expanded Measure of Net Worth," Working Paper. (July 2021)

Love, David A., Paul A. Smith, and Lucy C. McNair. 2008. "A New Look at the Wealth Adequacy of Older U.S. Households," Review of Income and Wealth, 54(4): 616-642.

Munnell, Alicia H., Anqi Chen, and Robert Siliciano. 2021. *The National Retirement Risk Index: An Update from the 2019 SCF,* Center for Retirement Research at Boston College No. 21-2. (January)

Sabelhaus, John. 2019. "Household Portfolios and Retirement Behavior," Paper presented at the Stanford Institute for Economic and Policy Research (SIEPR) 2019 *Working Longer and Retirement* Conference, October 10 -11, 2019.

Sabelhaus, John, and Alice Henriques Volz. 2019. "Are Disappearing Employer Pensions Contributing to Rising Wealth Inequality?" FEDS Notes. Washington: Board of Governors of the Federal Reserve System, February 1, 2019, https://doi.org/10.17016/2380-7172.2308.

Sabelhaus, John, and Alice Henriques Volz. 2021. "Social Security Wealth, Inequality, and Lifecycle Savings," forthcoming in *Measuring and Understanding the Distribution and Intra/Inter-Generational Mobility of Income and Wealth*, National Bureau of Economic Research, Studies in Income and Wealth. Chicago: University of Chicago Press.

Saez, Emmanuel and Gabriel Zucman. 2016. "Wealth Inequality in the United States Since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal of Economics*, 131(2):519–578.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving "Optimally" for Retirement?" *Journal of Political Economy*, 114(4), 607-643.

Smith, Matthew, Owen Zidar, and Eric Zwick. 2019. "Top Wealth in the United States: New Estimates and Implications for Taxing the Rich," U.S. Treasury Department. Available at: https://scholar.princeton.edu/sites/default/files/zidar/files/szz_wealth_19_07_19.pdf