

Wealth Inequality and Retirement Preparedness: A Cross-Cohort Perspective

John Sabelhaus and Alice Henriques Volz

August 2022

PRC WP2022-21

Pension Research Council Working Paper

Pension Research Council

The Wharton School, University of Pennsylvania

3620 Locust Walk, 3302 SH-DH

Philadelphia, PA 19104-6302

Tel.: 215.573.3414 Fax: 215.573.3418

Email: prc@wharton.upenn.edu

<http://www.pensionresearchcouncil.org>

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. All findings, interpretations, and conclusions of this paper represent the views of the authors and not those of the Wharton School or the Pension Research Council. © 2022 Pension Research Council of the Wharton School of the University of Pennsylvania. All rights reserved.

Wealth Inequality and Retirement Preparedness: A Cross-Cohort Perspective

Abstract

High and rising US 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 Baby 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

Keywords: Baby Boomers, Gen X, relative wealth distribution, wealth inequality

John Sabelhaus

Adjunct Associate Research Professor, University of Michigan
Non-Resident Senior Fellow, Brookings Institution
1775 Massachusetts Ave., NW
Washington, DC 20036
Email: jsabelhaus@gmail.com

Alice Henriques Volz

Chief, Microeconomics Surveys Section
Research and Statistics
Board of Governors of the Federal Reserve System
20th Street and Constitution Ave. NW
Washington, DC 20551
Email: alice.h.volz@frb.gov

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 et al. 2016; Smith et al. 2021). 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 Volz, forthcoming).

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 straightforward 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 US families are inadequately 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). Nevertheless, comprehensive household wealth should also include 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 Volz 2019, forthcoming). Finally, the aggregate net present value of all future social security benefits less taxes (or social security wealth, SSW) for working age and retired individuals is substantial, adding another \$23.9 trillion to household wealth in 2019 (Sabelhaus and Volz, forthcoming).

Including DB pensions and SSW adds nearly 50 percent to marketable wealth overall, but the distributional effect of their inclusion is uneven. DB wealth and SSW, in particular, are more evenly distributed than the narrow concept ('base wealth')—marketable assets less liabilities—captured directly in the SCF. Adding DB wealth and SSW to the SCF base wealth concept lowers the top 1 percent and top 10 percent wealth shares substantially in the 1995 to 2019 survey waves, by 5 and 9 percentage points, respectively (Sabelhaus and Volz, forthcoming). Including DB wealth and SSW does not change the trends in top shares, however, because the increasing wealth concentration of base wealth has not been offset by either faster growth or decreasing concentration of DB wealth and SSW.

The comprehensive wealth concept that includes DB wealth and SSW can be measured over the life cycle and across cohorts using the SCF in a pseudo-panel framework. 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 Volz, forthcoming). A large majority of the growth in average comprehensive wealth between 1995 and 2016 occurred at older ages. Adjusting the wealth measure for expected social security funding shortfalls makes the age differentials even larger, especially for the bottom half of the wealth distribution. In fact, average comprehensive wealth is estimated to be lower in 2019 than it was in 1995 for younger individuals in the bottom 50 percent of their wealth distribution. These findings motivate our approach to studying retirement adequacy.

In this paper, we use the comprehensive wealth measures we developed for analyzing levels and trends in wealth concentration 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 built on that same principle. An alternate approach is to compare observed, individual wealth holdings against the predictions of a calibrated life-cycle model. 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 to avoid these potential pitfalls. Our cross-cohort approach to studying retirement adequacy is based on *relative* wealth measures: how the wealth distribution of a cohort compares to the cohorts older than 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. Nevertheless, given a reference point—for instance, 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.

To quantify differences across cohorts, we use two approaches to compare wealth distributions. The first approach is creating *relative rank distributions* that answer the question 'where would an individual of a given cohort be if their wealth were mapped into the distribution

of an earlier cohort?’ The second approach is *percentile point comparisons*, that answer the question ‘how do the wealth holdings differ for an individual at a given percentile of the wealth distribution across cohorts and ages?’¹

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. We use the term ‘relative rank gap’ to describe the distance, in percentile points, between the two cohorts’ distributions. When a younger cohort is ahead of an earlier cohort, the relative rank gap will be positive (so that a negative gap suggests a shortfall).² 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 percentile at a time allows us to drill down into the relevance of the various wealth components across the wealth distribution and how they change over the life cycle and time.

As an example, consider comparing Early Baby Boomers (1940s cohort) to the 1930s cohort, their reference cohort, at age 60. We observe that Early Boomers’ wealth is generally on par with or above the 1930s cohort across the full wealth distribution. That is, the relative ranks of the 1940s cohort are all greater than the 1930s cohort ranks. As a result, when we look at the percentile point comparisons, the 1940s cohort’s comprehensive wealth value above the 1930s cohort’s comprehensive wealth value at the 10th, 25th, 50th, 75th, and 90th percentiles. What the percentile point comparisons point to is the key contribution of social security to that result, that the 1940s are ahead of 1930s, especially at the bottom of the wealth distribution where base wealth values are lower for the 1940s cohort.

Our dataset spans 30 years, and thus we can also compare the 1950s cohort to the 1930s cohort, still the reference cohort, at age 60. The relative wealth holdings of the 1950s cohort (the ‘Mid-Boomers’) at age 60 gives us the first indications of deterioration at the bottom of the wealth distribution. The extent of the shortfall, though, depends on the wealth concept. For example, the base concept (i.e., excluding SSW) shows a relative rank gap of between -5 and -10 for individuals in the 1950s cohort relative to the 1930s cohort in the bottom half of the wealth distribution. However, the comprehensive wealth measure including SSW completely reverses those gaps to the point where the rank gap is at least two for nearly all of the bottom half of the distribution.

Looking at younger ages, there is more evidence of relative deterioration in the bottom and middle of the wealth distribution for ‘Late Boomers’ (born in the 1960s) and ‘Gen-Xers’ (born in the 1970s). Base wealth at the 10th and 25th percentiles of the late Boomer and Gen-X distributions is well below the wealth of earlier cohorts observed at the same ages. In contrast, further up the wealth distribution, there are notable, positive relative rank gaps.

Social security is an important offset to base wealth declines at the bottom of the wealth distribution for younger cohorts. Adding SSW and estimating the comprehensive wealth distribution 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 for younger cohorts, with 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 coming from social security.

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 distributions, and the fifth shows percentile point comparisons. Section six concludes.

Data and Methods

Our approach to studying retirement adequacy requires a comprehensive, individual-level wealth measure across birth cohorts and over the life cycle. Achieving this goal involves starting with high-quality, household-level balance sheet data and adding in household-level estimates of SSW and DB wealth.³ The micro data used here are the Survey of Consumer Finances (SCF) for 1989 through 2019. The SCF is focused on household balance sheets, and the survey also has extensive information about incomes, demographics, and labor force experiences.⁴ The SCF alone is not ideal for calculating a comprehensive wealth measure; we bring in additional information on earnings and relative mortality to estimate DB and SSW more precisely than would be possible with the SCF alone. We estimate life-cycle earnings for individuals and their spouse/partners, required since the SCF is a series of cross-section snapshots, as inputs to calculating social security taxes and benefits over the life cycle. Since individual mortality rates are necessary to estimate present discounted values for both SSW and DB pensions, we differentiate mortality by age, sex, income, and birth year.

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 Volz, forthcoming). 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 less real-world relevance. In contrast, expected SSW is the present discounted value of total benefits less future taxes, conditioned on the individual working and paying taxes through a given expected retirement 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 life-cycle earnings patterns derived from linked longitudinal survey and administrative data in the Health and Retirement Study (HRS). The detailed assumptions needed to create the required inputs for computing expected SSW in the SCF are discussed in earlier papers (Sabelhaus 2019; Sabelhaus and Volz, forthcoming).

The measure of DB wealth consistent with the 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 modules. The three categories cover: (1) DB benefits already being received, (2) DB benefits associated with a past job where the known benefit amount will be received at a specific future date, and (3) 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. These benefit payments are the primary input to a present discounted value calculation that also involves an assumed interest rate and demographically differentiated mortality

rates (described below). This present discounted value calculation is equivalent to the level of financial assets that the retirement plan sponsor must hold to pay those promised benefits, and thus corresponds directly to the base household wealth measure captured in the survey.

Calculating DB pension wealth for workers covered by a plan on their current job is more complicated, because the benefit that will be received in the future is unknown as of the survey date. The aggregate measure of DB wealth that corresponds to our comprehensive household wealth is the financial liability of retirement plan sponsors; it corresponds to only 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, or accrued benefit obligation. 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 the differential mortality adjustment, which is key for estimating present values for retirement income streams.⁵ Constructing the present value of social security and DB pension incomes requires survival probabilities, which are computed for each SCF respondent and their spouse/partner, if present, through age 99. The starting point is Social Security Administration (SSA) cohort mortality by age and sex. These are modified by a differential mortality adjustment by income percentile within groups defined by age and sex based on Chetty et al. (2016). Chetty et al. (2016) map SSA death records onto income tax records to study mortality differences across the income distribution, which we transform into a differential adjustment (for methodological detail, see the appendix to Sabelhaus and Volz, forthcoming).

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 approximately 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 evaluation of retirement adequacy creates an index based on the SCF to estimate what fraction of the population is predicted to have enough resources (both wealth and income) at age 65 to maintain their pre-retirement standard of living, as measured by consumption (Munnell et al. 2021). The authors find that about half of current US 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 life-cycle model is an alternative way to measure retirement adequacy (e.g., Engen et al. 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 life cycle. Relatively simple life-cycle models can provide a wide range of predictions about optimal life-cycle 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 fall short of those average targets.

Studies defining ‘optimal’ wealth using structural life-cycle models that capture more heterogeneity across the population generally seem to suggest that fewer US families have inadequate retirement saving. One such study, using the Health and Retirement Study (HRS), makes use of detailed earnings histories and demographic characteristics and, thus, is able to fine tune life-cycle model predictions to specific types of households (Love et al. 2008). They find that only 18 percent of households in 2006 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 life-cycle optimization problem for every household in the HRS based on their unique characteristics (Scholz et al. 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 life-cycle model still depends on exactly which circumstances (or ‘state’ variables) are included in the model, and computational constraints limit which characteristics can be included. Furthermore, structural models are inherently static in nature and cannot 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 nor the static nature of structural life-cycle 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 generation relative to their parents, but changing

demographic and earnings patterns are leaving more individuals economically vulnerable (Butrica et al. 2007). The dynamic microsimulation draws attention to retirement wealth adequacy for specific groups whose circumstances are changing over time: 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 with a reduced-form, empirical approach (Brown et al. 2020). Their approach avoids lots of assumptions, and their findings further motivate our approach to studying retirement adequacy. The authors use early waves of the HRS, for whom economic and demographic characteristics are observed just before and through retirement, to assess which characteristics of pre-retirees are useful for predicting economic hardship during retirement. They use those correlations to predict how the cohorts approaching retirement today will fare in their retirement years. They find that those approaching retirement, particularly men, are indeed more likely to experience economic hardship, on average. Comparing cohorts at younger ages to help predict future outcomes is similar in spirit to what we do here. That is, what can we say about the likely outcomes for a cohort of pre-retirees 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 how well US households are prepared for retirement. One can look at the same individuals 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 investigation focuses on how the

wealth of younger cohorts compare to current retirees (or near-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 at a specific rank in the wealth distribution 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, 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 wealth do individuals hold at a given percentile of the wealth distribution across cohorts and ages?’ We call these charts *percentile point comparisons*. The percentile points comparison charts focus on a single slice of the wealth distribution across cohort and age, instead of characterizing the entire distribution like the relative rank charts. Another way to think of this is that the relative rank charts show us counts of people at points in the wealth distribution, while the percentile point comparisons are about dollar gaps for the people at one of those points. Though we can only focus on one wealth percentile at a time in a percentile point comparison, the data disentangle the contributions of different wealth sources across the distribution.

Relative Rank Distributions

We refer to *relative rank distributions* as the first approach we use to look at relative wealth distributions. . Our charts answer the question ‘what wealth rank would individuals of one cohort represent, if their wealth were mapped into the distribution of an earlier cohort?’ The comparisons presented are specific to an age group and a wealth concept, so there are several permutations such relative rank calculations.

We limit the number of comparisons by working with two wealth concepts across three age groups. As our base wealth measure, we use SCF net worth plus DB benefits, while comprehensive wealth—the base wealth measure plus SSW—is the alternative. The three age groups are ‘40’ (age 38-42), ‘50’ (age 48-52), and ‘60’ (age 58-62), with reference cohorts 1930s, 1940s, and 1950s, respectively, for these age groups. Given the SCF time span, the age 60 relative rank comparisons (for example) use the 1930s as the reference cohort, and the 1940s and 1950s cohorts as the two comparison groups.⁶ The age 50 and age 40 relative ranks shift the reference and comparison cohorts forward one and two decades, respectively.

The goal of the relative rank and percentile point comparisons is to avoid making absolute statements about what ‘adequate’ retirement wealth means, though the approach does require an assumption about what it means for wealth levels to be ‘comparable’ across cohorts. The reference and comparison cohorts are either 10 or 20 years apart in time, so comparing nominal wealth is misleading. To convert wealth to real dollars, the inflation adjustment uses the data series used to calculate and inflation-adjust social security benefits (the CPI).

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 if the comparison cohort relative rank distribution falls along the 45-degree line, the comparison cohort can be said to have the same wealth distribution as the reference cohort. If the relative rank distribution were above the 45-degree line, the comparison cohort wealth is above reference cohort wealth at that percentile.

With those chart-reading principles in mind, the first relative rank distribution in Figure 1a uses the base wealth concept (i.e., SCF net worth plus DB wealth) at age 60. The reference cohort for age 60 is the 1930s. The grey lines show the relative rank distributions for the 1940s (loosely, Early Boomer) cohort, and the dashed lines show the relative rank distributions for the 1950s (loosely, Mid-Boomer). The 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. The horizontal distance between the relative rank and 45-degree lines measures the relative rank gap. For example, an individual around the 55th percentile of the 1940s cohort had the same wealth as an individual at the 70th percentile of the reference cohort, for a relative rank gap of 15. In the bottom 30 percent of the distribution, the relative rank gap was never above five. But the rank gap grew to about 10 at the median, reaching the highest value at the 70th percentile. The relative rank distributions only show people, not dollars. Accordingly one cannot compare differences in the dollar amounts that correspond to the relative rank gaps, which we discuss that in the next section.

Figure 1 here

Although the relative rank distributions show an unambiguous improvement between the age 60 1930s and 1940s wealth distributions, the story is different for the 1950s cohort. The latter group tracks the first few percentiles of the 1930s cohort, but then it falls behind by a few rank

points for much of the distribution. Around 45 percent of the 1950s cohort had less wealth than their counterparts in the 1930s group. The fact that the relative rank gap is positive for the 1950s cohort at higher percentiles is consistent with what we know about rising wealth inequality over time. Again, all we can say is that an individual in the 1950s cohort at (say) the 25th percentile of the wealth distribution had the same wealth as someone in the 1930s at the 20th percentile, which seems like a modest gap.

Adding expected SSW to the base wealth concept pushes the 1940s cohort even further above the 1930s, and it also improves the relative ranks for the 1950s cohort (Figure 1b).⁷ The relative rank gap for the 1940s was above 10 for the 30th percentile of the reference cohort distribution, and the relative rank gaps were now all positive for the 1950s cohort. One way to interpret this shift in relative ranks is that social security was *relatively* more important for the 1940s and 1950s cohorts than it was for the 1930s reference cohort. Our approach to estimate SSW in the micro data captures the fact that changing demographics and lifetime wages produced differences in SSW for individuals at the same point in the wealth distribution across cohorts (Sabelhaus and Volz, forthcoming). If average individuals in the comparison cohort at a given age and wealth percentile had a longer life expectancy, they would be allocated more SSW. Changes in social security replacement rates, marriage patterns, relative earnings between spouses (through spousal and survivor benefits), and labor force participation also matter, but the relative ranks show that, on net, SSW was relatively more important for the younger cohorts.

Shifting back to the base wealth concept and looking earlier in the life cycle, the relative ranks at age 50 show that some of the 1950s cohort trailed the Early Boomer counterparts at low to modest wealth levels (Figure 2a). By contrast, the 1960s cohort trailed Early Boomers by a wide margin, approximately 10 rank points over much of the distribution. Individuals in the 20th -25th

percentiles of the 1960s cohort had only as much wealth as those between the 12th and the 15th percentile of the 1940s wealth distribution. The highest wealth individuals in both the 1950s and 1960s cohorts caught up to and surpassed their counterparts in the reference cohort, again consistent with rising wealth inequality over time. Yet, more of the 1950s cohort was relatively higher than the reference cohort, than was the 1960s cohort.

Figure 2 here

Adding SSW to the base wealth concept notably changes the relative rank gaps at age 50 (Figure 2b). Since both the reference cohort and the comparison cohort wealth distribution included SSW, the relative rank lines 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 implied relative improvement, compared to the base concept shown above (Figure 2a), as the 1950s now were ahead of the 1940s cohort and the relative rank gaps for the 1960s were notably smaller. For the 1960s cohort, there was a modest deterioration in retirement preparation using this framework (a relative rank gap of a few points at most), though the wedge appears for most of the wealth distribution.

As noted in the introduction, our interest in relative wealth adequacy across cohorts at various ages is partly motivated by the observation that average wealth at younger ages has fallen over time, which is another way of saying that a disproportionate share of the wealth gains accumulated to older age groups (Sabelhaus and Volz, forthcoming). Our approach allows us to consider these wealth declines at younger ages along two additional dimensions: where people fell in the wealth distribution, and for which cohorts wealth fell. The relative rank distributions for base wealth at age 40 confirm relative wealth declines at low wealth levels for the 1970s cohort, but there was little difference between the 1950s and 1960s cohort at low wealth levels (Figure

3a). It is worth noting here that a significant portion of the age 40 base wealth distribution was close to zero, with the 25th percentile of the reference cohort having around \$25,000. This point becomes clearer when we focus on the low wealth percentiles in the percentile point comparison charts. Further up the distribution, the 1960s cohort did relatively better than the 1950s cohort, with a rank gap around 10 between the 60th and 80th percentiles of the reference cohort.

Figure 3 here

Social security was shown to be relatively important for relative rank distributions at age 50, which also held at age 40, such that the bottom quarter of the distribution for both the 1960s and 1970s cohorts had positive rank gaps (Figure 3b). It is particularly notable that including SSW closed the relative rank gaps across the whole wealth distribution for the 1970s cohort. Using a comprehensive wealth measure that includes scheduled SSW—payable or not—the 1960s and 1970s cohorts were slightly above the 1950s cohort.

Tying together the relative rank distributions at various ages using the most comprehensive wealth measure (Figures 1b, 2b, and 3b), one can assert that (1) the 1950s cohort tracked or was slightly behind the 1930s cohort at age 60; (2) the 1970s cohort tracked or was slightly behind the 1950s cohort at age 40; and thus, if the relative ranks are transitive across time, (3) the 1970s cohort was somewhere between tracking and slightly behind the 1930s cohort.

The optimistic conclusions from the relative ranks at age 40 are contingent on whether the increasing SSW offsets relative declines in other types of wealth are in fact payable. Our assumption used to estimate payable benefits is simple: we adopt the SSA actuaries' projection that 78 percent of benefits are payable after 2033. Thus, payable SSW is the present value of scheduled benefits through 2033 plus 78 percent of scheduled benefits after. The timing is such that the youngest cohorts will face the largest lifetime cuts, and this result comes through in the

relative rank distributions. When we compute relative ranks at age 40 using only payable SSW, there is more evidence of relative shortfalls for the 1960s and 1970s cohorts across much of the wealth distribution. The 1950s reference cohort will also face benefit cuts in the payable scenario, albeit proportionally less. The largest decrease in the rank gap for the 1960s cohort is around the 25th percentile of the reference distribution, while the 1970s cohort now experiences notable rank gaps across the bottom two-thirds of the distribution (see Figure 3c). To evaluate how relative wealth declines under the payable benefits scenario, we examine this question next.

Percentile Point Comparisons

Cross-cohort shifts in relative rank positions are useful because they help us compare an entire cohort group relative to another. The shortcoming is that we cannot assess how large any shortfalls are because the charts are ordinal in nature. Accordingly, we next present *percentile point comparisons* 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 of this approach 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 choice. Although it is obvious what it means to be at a given percentile for a given wealth concept, comparing 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 and then add average payable SSW and average scheduled but not payable SSW. The results are comparable across these two approaches but show slightly lower average base wealth compared to if one were to rank individuals by comprehensive wealth instead and calculate the components.

To build the percentile points charts, we start with comprehensive wealth (SCF net worth plus DB and scheduled SSW) for a given percentile-cohort-age combination and solve for the 10th percentile of that wealth measure. We then repeat the process for the other two measures: SCF net worth plus DB and payable SSW and the base wealth measure. The 10th percentile of SCF net worth plus DB and payable SSW is then solved for as the difference between the first two values, and the 10th percentile of SCF net worth plus DB is the difference between the second and third values.

We begin the percentile comparisons near the bottom of the wealth distribution, at the 10th percentile (Figure 4). The differences in overall bar heights within a given age and cohort comparison are dollar-valued analogous to the relative rank gaps at the 10th percentile using the comprehensive wealth measure. The fact that overall bar heights are generally similar is the parallel to the fact that the relative ranks were generally close to the 45-degree line; that is, the relative rank gaps were near zero, at the 10th percentile in the comprehensive wealth charts. One exception—the 1940s cohort at age 60 had more wealth at the 10th percentile than did the 1930s cohort—is analogous to the observation that the 1940s relative rank line was above the 45-degree line at low wealth levels (Figure 1b). 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 was dramatically lower than predecessor cohorts based on payable benefits alone.

Figure 4 here

At the 25th percentile of the wealth distribution, base wealth plays a slightly more important role in household resources, but social security still dominates (Figure 5). The relative deterioration in comparison cohorts' base wealth at the 25th percentile is notable for the 1970s cohort at age 40, the 1960s cohort at age 50, and the 1950s cohort at age 60. These shortfalls reflect the relative rank gaps of around -10 for the base wealth concept (Figures 1a, 2a, and 3a) at the 25th percentile. In contrast, the 1940s cohort at age 60 had higher base wealth. As with the 10th percentile, SSW reverses some of the relative wealth decline at the 25th percentile, but focusing on payable benefits alone, the prospects for low wealth individuals in the 1960s and 1970s cohorts are bleaker.

Figure 5 here

The contributions to total wealth from social security and non-social security sources are more balanced in the middle of the wealth distribution (Figure 6). Similar observations about relative declines in base wealth present at the 25th percentile also show up at the 50th percentile, but now only for the 1970s cohort at age 40 and the 1960s cohort at age 50, analogous to the negative relative rank gaps near the middle of the wealth distribution in those two instances. At the median, the 1950s cohort is comparable to the reference cohort at age 60. Nevertheless, the relative changes are more muted or are a relative improvement over the reference cohort using the comprehensive wealth measure, consistent with nearly all relative ranks above the 45-degree line at the 50th percentile (Figures 1b, 2b, and 3b). In general, the 50th percentile comprehensive wealth measures show less relative deterioration across cohorts when compared to lower wealth levels.

Figure 6 here

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 7). In all cases, the data show gains in relative comprehensive wealth at the 75th percentile, although focusing on only payable SSW leads to some small relative shortfalls. 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, at least compared to the bottom half of the wealth distribution, is that social security solvency is a relatively less important determinant of the prospects for currently young, higher wealth individuals.

Figure 7 here

For a complete presentation of the wealth distribution, we also show wealth components at the 90th percentiles (Figure 8). Here there are few concerns for retirement adequacy: the impression is the mirror image of the 10th percentile charts, in which SSW dominates. At the 90th percentile, social security is a relatively small component of comprehensive wealth, and the issue whether social security is payable or not has very modest implications. At the 90th wealth percentile, slightly different than the takeaway from the 75th percentile—younger cohorts do look better off. The 90th percentile point comparisons are also consistent with the idea that rising wealth inequality is more about what's happening within the top 10 percent, and less about the 90th percentile relative to other groups per se, since there are only modest improvements of the comparison cohorts relative to the reference cohort.

Figure 8 here

Conclusions

Are the 60-year-olds who were born in 1960 well-prepared for retirement? Are the 50-year-olds born in 1970, and the 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? Researchers with different retirement wealth adequacy yardsticks could 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 at different 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 wealth declines (and presumably retirement preparedness) in the bottom half of the distribution.

The *relative* wealth approach means we cannot say anything about absolute retirement preparedness for a given individual or a cohort. Yet if we believe that 20 percent of the 1930s cohort experienced financial hardship in retirement, which seems probable given Brown et al.'s (2020) analysis of the HRS, then some fraction below 20 percent of the 1940s cohort can be expected to experience financial hardship in retirement, based on their relative wealth distributions at age 60.⁸ Depending on the wealth concept used, we conclude that somewhere between the 13th and 18th percentile of the 1940s cohort had as much wealth as the 20th percentile of the 1930s

cohort at age 60, putting them moderately or significantly ahead of the 1930s cohort depending on whether social security is considered.

Also, and still depending on the specific wealth measure 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, the 30th percentile of the 1970s cohort had the same wealth as the 20th percentile of the 1950s cohort at age 40, meaning we could be looking at a 50 percent increase in the fraction of retirees facing hardship when the 20th percentile is our cutoff, presuming that the 1950s cohort have a similar rate of hardship in retirement as did the 1930s cohort. At the bottom of the wealth distribution, relative retirement preparedness depends on how policymakers address expected future social security shortfalls.

References

- Bayer, P. and K. K. Charles (2018). 'Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940,' *Quarterly Journal of Economics*, 133(3): 1459-1501.
- Bhutta, N., J. Bricker, A. C. Chang, L. J. Dettling, S. Goodman, J. W. Hsu, K. B. Moore, S. Reber, A. H. Volz, and R. 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, J., A. Henriques, J. Krimmel, and J. Sabelhaus (2016). 'Measuring Income and Wealth at the Top Using Administrative and Survey Data.' *Brookings Papers on Economic Activity*, 2016(1): 261-321.
- Brown, J., K. Dynan, and T. Figinski (2020). 'The Risk of Financial Hardship in Retirement: A Cohort Analysis.' In O. S. Mitchell and A. Lusardi, eds., *Remaking Retirement: Debt in an Aging Economy*. Oxford, UK: Oxford University Press, pp. 60-85.
- Butrica, B., H. M. Iams, and K. E. Smith (2007). 'Understanding Baby Boomers' Retirement Prospects.' In B. Madrian, O. S. Mitchell, and B. J. Soldo, eds., *Redefining Retirement: How Will Boomers Fare?* Oxford, UK: Oxford University Press, pp. 70-91.
- Engen, E. M., W. G. Gale, and C. E. Uccello (1999). 'The Adequacy of Household Saving.' *Brookings Papers on Economic Activity*, 1999(1): 65-187.
- Hurd, M. J. and S. Rohwedder (2012). 'Economic Preparation for Retirement.' In D. Wise, ed., *Investigations in the Economics of Aging*. Chicago, IL: University of Chicago Press, pp. 77-117.

- Love, D. A., P. A. Smith, and L. 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, A. H., A. Chen, and R. Siliciano (2021). 'The National Retirement Risk Index: An Update from the 2019 SCF.' Center for Retirement Research at Boston College No. 21-2. Boston, MA: Center for Retirement Research at Boston College.
- Sabelhaus, J. (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, J. and A. H. Volz (2019). 'Are Disappearing Employer Pensions Contributing to Rising Wealth Inequality?' *FEDS Notes*. February 1: <https://doi.org/10.17016/2380-7172.2308>.
- Sabelhaus, J. and A. H. Volz (forthcoming). 'Social Security Wealth, Inequality, and Lifecycle Saving.' In R. Chetty, J. N. Friedman, J. C. Gornick, B. Johnson, and A. Kennickell, eds., *Measuring Distribution and Mobility of Income and Wealth*. Chicago, IL: University of Chicago Press.
- Saez, E., and G. 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, J. K., A. Seshadri, and S. Khitatrakun (2006). 'Are Americans Saving "Optimally" for Retirement?' *Journal of Political Economy*, 114(4): 607-643.
- Smith, M., O. Zidar, and E. Zwick. (2021). 'Top Wealth in America: New Estimates and Implications for Taxing the Rich.' NBER Working Paper No. 29374. Cambridge, MA: National Bureau of Economic Research.

Endnotes

¹ These are analogous to the ‘rank gap’ and ‘level gap’ concepts described in Bayer and Charles (2018) to discuss the evolution of the racial earnings gap. Our framework is slightly different in that we are comparing cohorts at different ages (instead of indexing to a calendar year), and we present the full distribution of relative ranks.

² This is the opposite presentation of the rank gap as used by Bayer and Charles (2018).

³ A more detailed explanation of the methods used to construct comprehensive household wealth can be found in Sabelhaus and Volz (2021).

⁴ See Bhutta et al. (2020) for a discussion of the SCF and the most recent results, the 2019 survey.

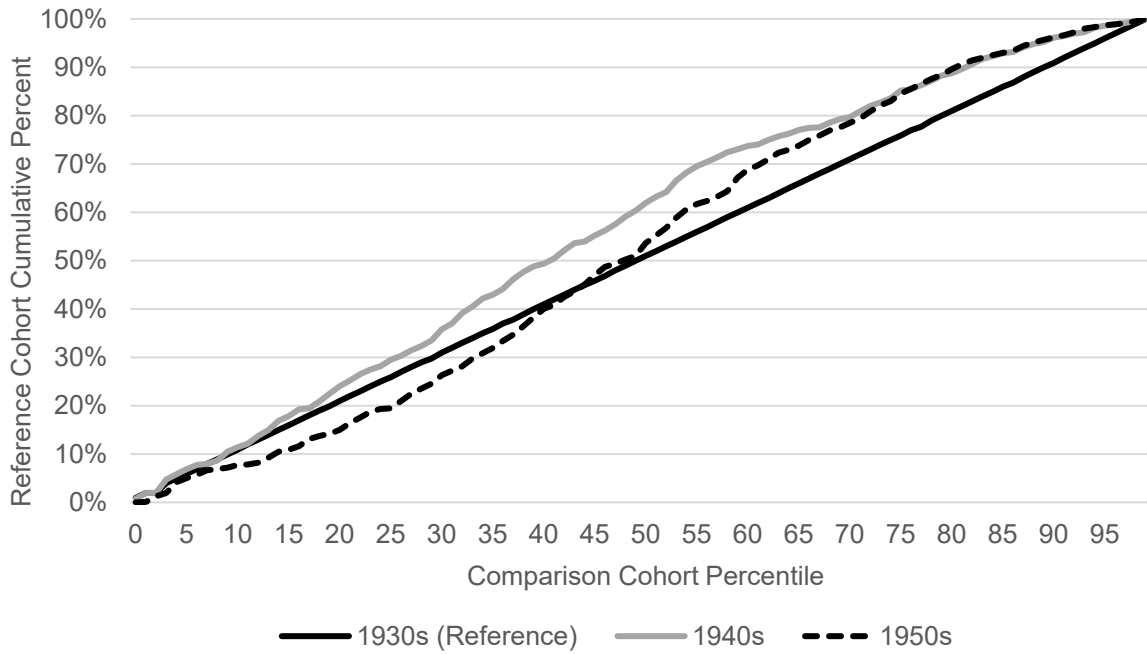
⁵ The demographic variables are available for both the SCF respondent and the spouse/partner, which is important, because the SCF micro files always assign the male in a couple (person closest to age 40 in a same sex couple) to be the respondent, and the other individual to be the spouse/partner, without regard to which is the primary earner.

⁶ Technically, because the SCF data spans 1989 to 2019, we use the last seven years of each decade to represent a birth cohort, so the ‘1930s’ means 1930-37 and ‘1970s’ is 1970-77. 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, that is, each birth cohort gets equal weight.

⁷ There is only minor reranking of individuals when changing from the base wealth distribution to the comprehensive wealth distribution, as most households expect to receive Social Security benefits and the level of SSW is highly correlated with base wealth levels. Thus, an individual’s rank is very similar across Figures 1a and 1b.

⁸ Brown et al. (2020) measure hardship of the 1932-37 birth cohort in 2014, when they were ages 77-82. They find that 6.5% of individuals were in poverty while 12.4% had annuitized wealth below 150% of the poverty line. Since this is a snapshot in time, it is likely that some fraction of those not *currently* in hardship experienced some difficulties earlier in retirement.

1a. Base wealth (net worth plus DB)



1b. Comprehensive wealth (net worth plus DB and SSW)





Figure 1. Relative rank distributions at age 60

Source: Authors' calculations.

2a. Net worth plus DB wealth



2b. Net worth plus DB and SSW






Figure 2. Relative rank distributions at age 50

Source: Authors' calculations.

3a. Base wealth (net worth plus DB)



3b. Comprehensive wealth (net worth plus DB and SSW)



3c. Net worth plus DB and scheduled and payable SSW




Figure 3. Relative rank distributions at age 40

Source: Authors' calculations.




Figure 4. 10th percentile of cross-cohort wealth distributions

Source: Authors' calculations.




Figure 5. 25th percentile of cross-cohort wealth distributions

Source: Authors' calculations.





Figure 6. 50th percentile of cross-cohort wealth distributions

Source: Authors' calculations.



Net Worth Plus DB Benefits
 Payable Social Security Benefits
 Scheduled but not Payable Social Security Benefits

Figure 7. 75th percentile of cross-cohort wealth distributions

Source: Authors' calculations.

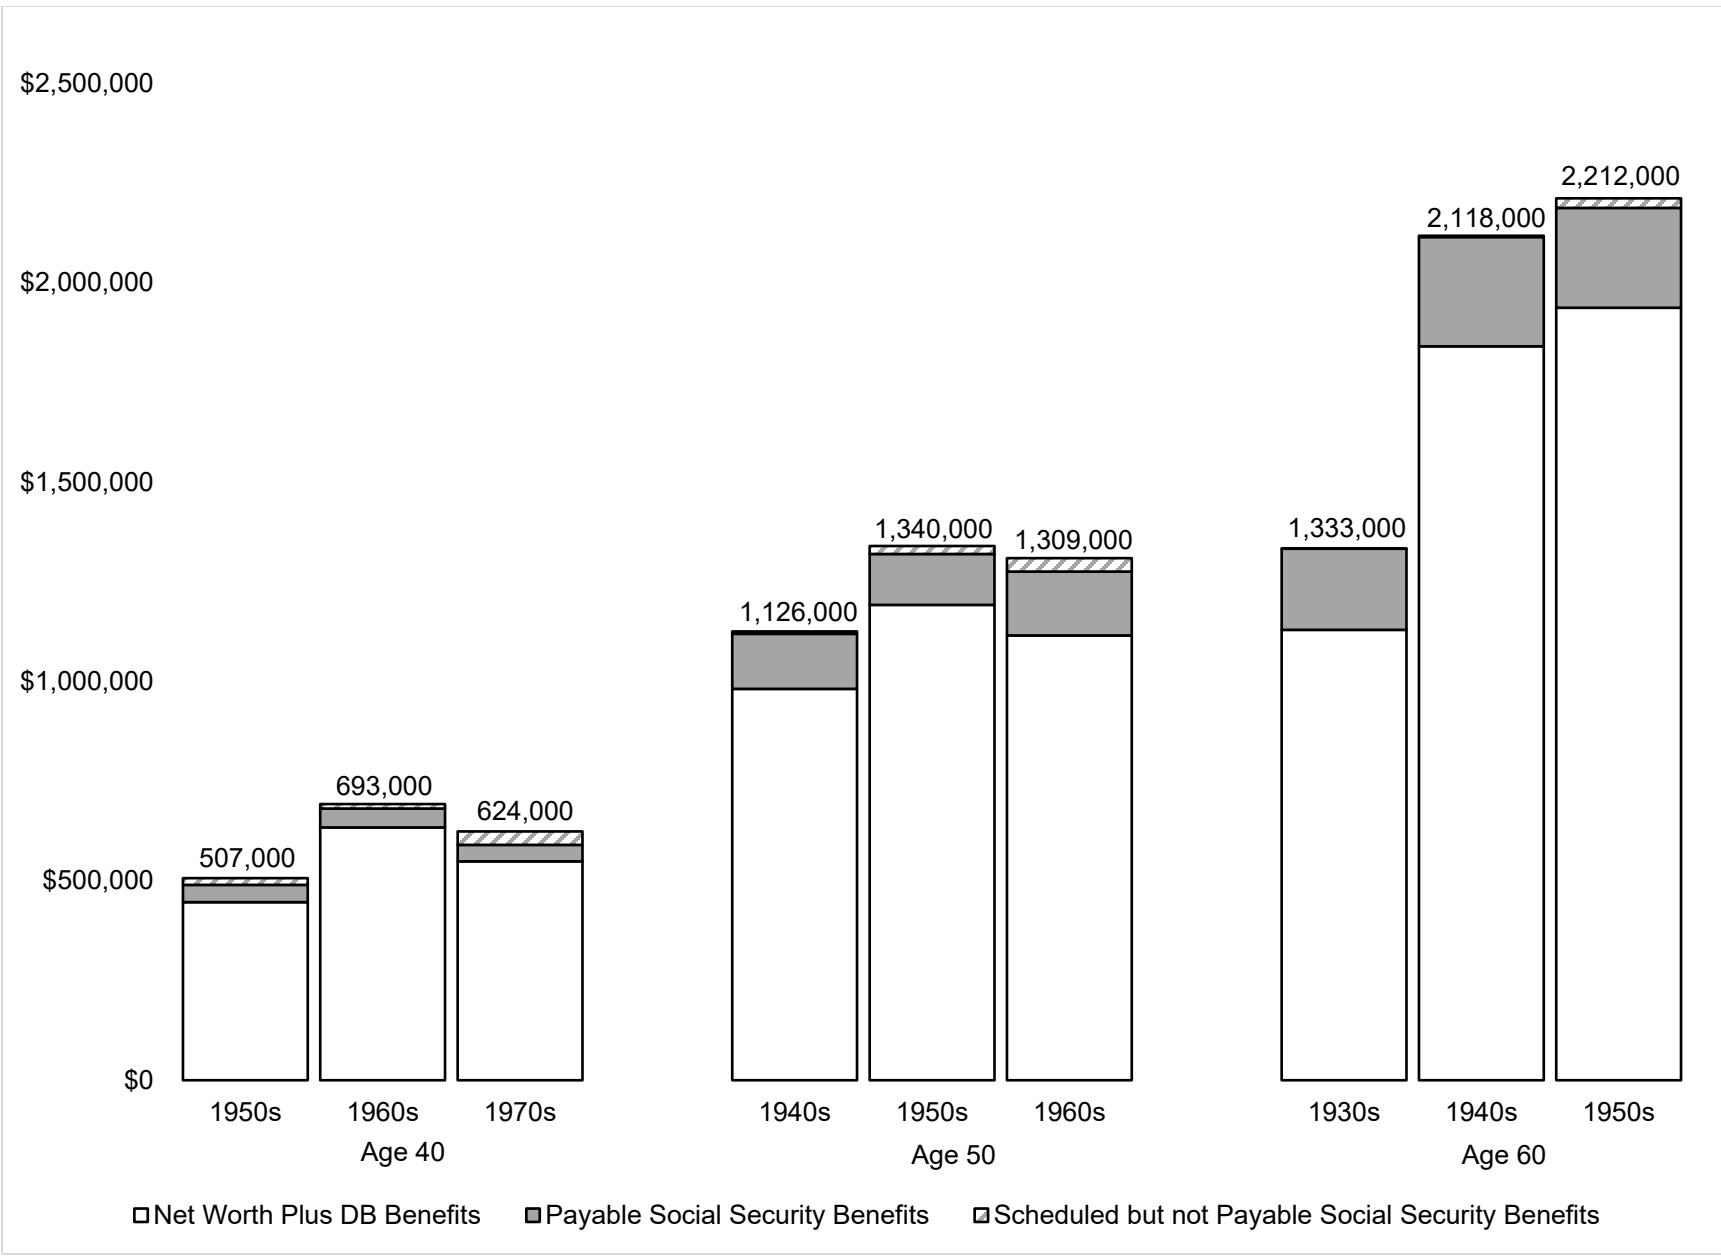


Figure 8. 90th percentile of cross-cohort wealth distributions

Source: Authors' calculations.