



Social Security Wealth, Inequality, and Life-cycle Saving: An Update

John Sabelhaus and Alice Henriques Volz

MRDRC WP 2020-416

UM20-16

Social Security Wealth, Inequality, and Life-cycle Saving: An Update

John Sabelhaus
University of Michigan

Alice Henriques Volz
Federal Reserve Board

November 2020

Michigan Retirement and Disability Research Center, University of Michigan, P.O. Box 1248.
Ann Arbor, MI 48104, mrdrc.isr.umich.edu, (734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium through the University of Michigan Retirement and Disability Research Center Award RDR18000002-02. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA, the Board of Governors of the Federal Reserve System or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Regents of the University of Michigan

Jordan B. Acker; Huntington Woods; Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Paul W. Brown, Ann Arbor; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Ron Weiser, Ann Arbor; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*



Social Security Wealth, Inequality, and Life-cycle Saving: An Update

Abstract

Social Security wealth (SSW) is the present value of future benefits an individual will receive less the present value of future taxes they will pay. When an individual enters the labor force, they generally face a lifetime of taxes to pay before they will receive any benefits and, thus, their initial SSW is generally low or negative. As an individual works and pays into the system their SSW grows and generally peaks somewhere around typical Social Security benefit claiming ages. The accrual of SSW over the working life is most important for lower income workers because the progressive Social Security benefit formula means that taxes paid while working are associated with proportionally higher benefits in retirement. We estimate SSW for individuals in the Survey of Consumer Finances (SCF) for 1995 through 2019 using detailed labor force history and expectations modules. We use a pseudo-panel approach to empirically demonstrate life-cycle patterns of SSW accumulation and drawdown. We also show that including SSW in a comprehensive wealth measure generally reduces estimated levels of U.S. wealth inequality, but does not reverse the upward trend in top wealth shares.

Citation

Sabelhaus, John, and Alice Henriques Volz. 2020. "Social Security Wealth, Inequality, and Life-cycle Saving: An Update." Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper; MRDRC WP 2020-416.

<https://mrdrc.isr.umich.edu/publications/papers/pdf/wp416.pdf>



Notes on update

This working paper is an update of our National Bureau of Economic Research (NBER) Working Paper 27110, released in May 2020. That paper is forthcoming in the Conference on Research in Income and Wealth (CRIW) volume, *Measuring and Understanding the Distribution and Intra/Inter-Generational Mobility of Income and Wealth*. This update uses the additional year of Survey of Consumer Finances (SCF) data (survey year 2019) that became available in September 2020. There is also a methodological change with respect to differential mortality, and a new appendix that describes on-going research using Health and Retirement Study (HRS) detailed earnings histories to improve life-cycle earnings imputations in the SCF. Neither the update to 2019 nor changes in differential mortality methods affect the conclusions in the NBER version.

The new method for assigning mortality differentials by income is based on the work of Chetty et al. (2016), and replaces the approach we previously used based on the Congressional Budget Office long-term model and adapted to the SCF by Feiveson and Sabelhaus (2019). The Chetty et al. (2016) approach uses income tax records linked to Social Security mortality records, and the larger and more recent sample sharpens the correlation between income and mortality across the gender, age, and income distribution. Our previous approach assigned mortality differentials by three broad age groups and income quintile, while the Chetty et al. (2016) estimates make it possible to assign differentials by individual ages and income percentiles. See Appendix B for details about the improved differential mortality adjustments.

One of the key methodological contributions in this study is estimating detailed life-cycle earnings profiles for SCF respondents. Appendix C describes on-going research to improve those estimates using Social Security Administration (SSA) earnings records linked to HRS longitudinal records. Access to the linked HRS/SSA earnings data was delayed until near the end of the performance period for this grant, and thus Appendix C is mostly focused on the first steps in using HRS earnings profiles: reconciling SCF and HRS linking variables. The two data sets are first aligned in terms of basic demographics including birth year, gender, education, and occupation. The second step is to look at heterogeneity within types by life-cycle labor force attachment — number of years worked full and part time. Our on-going research is focused on how best to use observed earnings profiles for a given demographic type and pattern of labor force attachment to improve the imputations of life-cycle earnings in the SCF.

1. Introduction

Wealth inequality in the United States is high and has been increasing, measured using either the income-capitalization approach (Saez and Zucman 2016; Smith, Zidar, and Zwick 2019) or wealth surveys (Bricker, Henriques, Krimmel, and Sabelhaus 2016). However, these estimates of wealth inequality do not include the present value of Social Security benefits less the present value of Social Security taxes — or Social Security wealth (SSW). This omission is important because Social Security dominates other forms of retirement income for most families (Sabelhaus and Volz 2019) and Social Security has been growing relative to the size of the overall economy. A comprehensive measure of household wealth that includes SSW provides new insights into overall wealth inequality and our understanding of life-cycle saving behavior.

Aggregate Social Security benefits in the U.S. are now roughly the same order of magnitude as the sum of total pension benefits received and retirement account withdrawals. Social Security benefits are also growing relative to the economy's size, from roughly 4% of personal income in the early 1970s to almost 6% today. Focusing on just the benefit side of SSW, the present value of Social Security benefits for everyone who has paid anything into the system was \$73.3 trillion in 2019.¹ Thus, the present value of Social Security benefits is estimated to be roughly double all other household-sector pension and retirement account assets combined, and approximately three-fourths the size of all conventionally measured household net worth. Social Security is

¹ Table VI.F2 in the 2019 Trustees Report, Office of the Chief Actuary, Social Security Administration, available at https://www.ssa.gov/OACT/TR/2019/VI_F_infinite.html#1000308.

also an important retirement wealth equalizer, as employer-sponsored pension and retirement accounts accrue disproportionately to high wealth families (Sabelhaus and Volz 2019).

Unlike pensions and other forms of retirement wealth, aggregate net SSW is zero, because there are no associated private claims to future production reflected in the market value of financial assets. The relatively small OASDI trust fund in principle represents real claims of Social Security beneficiaries on the rest of government, but even those claims are in the form of government bonds that are also the liabilities of taxpayers. Individual SSW (the present value of future benefits less the present value of future taxes for a given person) does vary over the life cycle, meaning some (usually older) participants have positive SSW, while others (the young) have negative SSW. Indeed, the rise and fall of SSW over the life cycle is an important form of saving and dissaving that is measurable, just like any other type of life-cycle wealth.

SSW also varies within age groups by lifetime income because of the progressive Social Security benefit formula, and thus lower lifetime income workers have higher SSW accrual rates (relative to income) while they are working. In a life-cycle sense, workers with low lifetime earnings are “saving” much more than higher lifetime income groups at any given working age, because each tax dollar paid is associated with a larger increase in the net present value of benefits they will ultimately receive. The prediction for conventionally measured saving and household wealth is clear. Lower lifetime earners with relatively high SSW can achieve the same total wealth (relative to income) with less (conventionally measured) saving.

The expected interplay between conventionally measured household saving and SSW is by no means a new idea, but the empirical evidence is limited by the lack of a comprehensive data set with both types of wealth for the same individuals. In this paper, we begin with the triennial Survey of Consumer Finances (SCF), which has high-quality income and balance sheet data for a representative sample (including a high-wealth oversample) of the population. We construct measures of SSW for individuals using the retrospective work history and prospective work expectations modules in the SCF with imputed earnings based on the Health and Retirement Study (HRS), solving for future payroll taxes, own worker benefits, and spouse and survivor benefits. We then use a pseudo-panel methodology — tracking birth cohorts across the nine survey waves between 1995 and 2019 — to study life-cycle SSW patterns and how SSW interacts with other balance sheet components to determine overall wealth inequality.

There are four major takeaways from this exercise. First, aggregate SSW is quantitatively important when compared to other components of household wealth. As we detail in the next section, there are multiple conceptual issues when estimating SSW, and the choice of discount rate and counting scheduled versus payable benefits has a big impact on any given SSW empirical estimate. Using what we refer to below as “expected scheduled” SSW and our preferred 2.8% real discount rate, we estimate that the present discounted value (PDV) of future benefits less future taxes for all SCF respondents and their spouse/partners in 2019 was about \$24 trillion, the difference between a PDV of benefits of \$38 trillion and a PDV of taxes of \$14 trillion. However, we estimate that only \$19 trillion of the \$24 trillion in SSW is payable under current law. For some perspective, the published measure of household net worth in the SCF was

around \$96 trillion in 2019 (Bhutta et al. 2020) and the PDV of Defined Benefit (DB) pension benefits was \$19 trillion (Sabelhaus and Volz 2019). Our estimate of SSW in 2019 is higher — \$33 trillion — using Office of the Chief Actuary (OCACT) discount rates.²

The second main finding is that SSW — unsurprisingly — is quantitatively more important for otherwise low-wealth families at any given age. There is a myriad of ways to measure wealth inequality, and although the alternative approaches to sorting or constructing summary measures affect estimated levels of inequality at any point in time, there are common trends across the various measures, so little is lost by focusing on one specific distributional approach. Our baseline wealth inequality measure is the sum of published SCF net worth and our estimated household-level DB wealth (Sabelhaus and Volz 2019), what we refer to as “household” wealth. Our baseline sorting mechanism involves assigning a given percentage of SCF respondents and spouses/partners (not households) within each age group into the corresponding wealth groups, which eliminates confounding life-cycle effects (wealth generally increases with age) on wealth grouping. Thus, our “top 10%” includes the top 10% of respondents and

² The conceptually equivalent OCACT published SSW is \$38 trillion for 2019. The SCF-based aggregate SSW estimate is only slightly less than OCACT, but that obscures larger differences in both sides of the SSW computation (the PDV of benefits and the PDV of taxes). As discussed later in the paper, we are missing future benefits (and taxes) for persons 15 and older who are not SCF respondents or spouses/partners. We don’t simulate future transitions onto disability, and we don’t count spouse/survivor benefits because of future or past marriages. The omission of disability may have a substantial effect on our SSW estimates because individuals have a 25 percent probability of becoming disabled before age 65. See, for example, the OCACT annual actuarial note on disability, available at <https://www.ssa.gov/oact/NOTES/ran6/an2020-6.pdf>.

spouses/partners within each age group, as ranked by the sum of published SCF wealth and our estimated DB wealth.

Sorting by household wealth within age groups and person-weighting implies that our measures always answer questions of the form, “how does expanding the wealth concept affect the average wealth holdings at a given age of a conventionally measured wealth group?” We show that SSW is relatively much more important for low-wealth individuals. For example, the bottom 50% of persons ages 35 to 44 in 2019 had average household wealth around \$22,000. However, the same group had average expected SSW of just over \$50,000, the difference between a PDV of benefits around \$137,000 and a PDV of taxes around \$87,000. This is unsurprising given that low-wealth individuals have much lower lifetime incomes, and the Social Security tax and benefit formulas are inherently progressive, even though differential mortality offsets some of that redistribution. In contrast, the top 10% of persons aged 35 to 44 in 2019 had, on average, about \$2,000,000 of household wealth. Their expected SSW was \$86,000, the difference between a PDV of benefits around \$254,000 and a PDV of taxes around \$169,000. Thus, although SSW rises with other wealth components (because both increase with lifetime income) the ratio of SSW to household wealth falls as household wealth rises.

The third major takeaway is that although incorporating SSW into household wealth has a substantial impact on wealth inequality *levels*, it does not change overall *trends* in top wealth shares. For example, while the top 10% share of household wealth (within age-sorted and person-weighted) increased from 53% to 63% between 1995 and 2019, the expanded wealth share that includes SSW increased from 45% to 55%. In an

important sense, adding SSW deepens, rather than ameliorates, concerns about rising wealth inequality, because roughly the same percentage point change in wealth share is applied to a much lower base. In addition, adding SSW does not reverse the exploding average wealth gap between the young and old, especially if we measure wealth using payable SSW.

The final takeaway is based on connecting the estimated SSW values across cross-section survey waves for 10-year birth cohorts. The triennial structure of the SCF gives us up to nine pseudo-panel SSW observations per birth cohort over the quarter-century between 1995 and 2019. By connecting the cohort averages between survey waves and drawing out the life-cycle patterns of SSW by age, we show how SSW starts out negative at young ages, increases steadily through retirement, and then gradually decreases as the remaining expected years of life (and years of expected benefit receipt) decline at older ages. Age ranges in which multiple cohorts overlap show us how average SSW is evolving across cohorts, due to a combination of demographic (through spouse/survivor benefits) and lifetime earnings effects.

The life-cycle shapes of SSW accumulation and decumulation are similar across wealth groups, but there are important differences in SSW relative to income that can be interpreted in terms of life-cycle saving rates. At age 60, average SSW relative to average permanent income is roughly 600% for the bottom 50% wealth group, 400% for the next 40% wealth group, and 100% for the top 10% wealth group.³ Given the observed average income growth between ages 30 and 60 across the three wealth groups, those ratios translate roughly into annual saving rates of (accumulation of SSW

³ See Bhutta, et al (2020) for a discussion of the SCF permanent income measure.

relative to income) of 15%, 12%, and 5% as we move from the lowest to highest wealth groups. In that sense, low- and middle-wealth families are saving a much larger fraction of their income than what is suggested by conventional saving measures (Feiveson and Sabelhaus 2019). One could say that low- and middle-wealth families just save differently than high wealth families, by participating in Social Security.

2. Social Security wealth

A given individual's Social Security wealth (SSW) is the present discounted value of their future benefits less the present discounted value of their future taxes. SSW estimates can be constructed at any point in the life cycle using earnings histories and various assumptions about continued work and benefit claim ages. The discussion here focuses on two main scenarios, which we refer to as the "expected" and "termination" approaches to measuring SSW as of a given age. The expected concept is based on respondent-reported future work expectations, and thus includes future taxes and benefits. The termination concept — borrowed from the Defined Benefit (DB) pension lexicon — assumes the participant does not work in the future, and thus, does not accrue additional benefits beyond what they have already earned. The two SSW measures in principle provide different perspectives on the evolution of SSW over the life cycle, and on how policy changes will affect the life-cycle profiles of SSW.

Social Security wealth defined

In the most generic sense, SSW for an individual at age a is the discounted present value of survival-adjusted benefits less taxes,

$$SSW_a = \sum_{s=a+1}^T (\text{benefits}_s - \text{taxes}_s) \Pi_{s/a} \beta^{s-a}$$

where T is the maximum lifespan, $\Pi_{s/a}$ is the probability of surviving to age s as of age a , and β is a period discount factor.

Moving from the generic representation to specific empirical measures requires introducing life-cycle timing for benefits and taxes, which in turn depend on the underlying life-cycle patterns of earnings and Social Security system parameters. Denote the last age that the individual works using m , and the first age that benefits are received using n . The specific benefit formula for an individual depends on their birth cohort, c . Benefits at age s are then given by,

$$\text{benefits}_s = b(\tilde{y}_m, n, c) \quad \forall s \geq n, 0 \text{ otherwise}$$

where \tilde{y}_m is the vector of earnings through age m , and $b(\cdot)$ is a benefit calculator that maps lifetime earnings through age m and benefit start age n into a benefit level at age s for a member of birthyear cohort c .

Taxes paid at age s depend on earnings at age s and the calendar-year specific tax rate τ and taxable maximum y^{\max} . Calendar year is the sum of cohort birthyear (c) and current age (s), so the tax parameters can also be written in terms of $c + s$. That is, for all ages s through the last working age m ,

$$\text{taxes}_s = \tau_{c+s} * \max(y_s, y_{c+s}^{\max}) \quad \forall s \leq m, 0 \text{ otherwise}$$

In the specific estimates described below, we will generally assume that benefits commence when working ends ($n = m+1$) or at the earliest possible age for benefits (generally 62) if the individual stops working prior to the earliest benefit age, but there is nothing in the notational framework above that requires those conditions.

When we adopt the simplifying assumption that benefit start age $n=\max(m+1,62)$, SSW is simplified to a two-dimensional object in current age (a) and stop work age (m).

For a given member of cohort c ,

$$SSW_a^m = \left[\sum_{s=\max(\max(a+1,m+1),62)}^T b(\tilde{y}_m, n, c) \Pi_{s/a} \beta^{s-a} \right] - \left[\sum_{s=a+1}^m \tau_{c+s}^* \max(y_s, y_{c+s}^{\max}) \Pi_{s/a} \beta^{s-a} \right]$$

Various combinations of a and m answer different questions about the role of SSW in life-cycle wealth accumulation and inequality. Measures of *expected* SSW are most useful for thinking about variation in money's worth or progressivity across and within cohorts, or overall system finances. Measures of *termination* (or conditional termination) SSW are more useful for capturing the incentive effects of over the life cycle.

Money's worth and system finances

One useful reference point for stop work age (m) is the *expected* value. When m is set to the expected (or already-realized for current beneficiaries) stop work age values (denoted by $*$), the average across all members of a given birth cohort of the constructed SSW_a^* map out the mean life-cycle pattern of SSW by age for that cohort. Those typical life-cycle SSW patterns will vary across birth cohorts, depending on the Social Security tax and benefit parameters faced by that cohort, their lifetime earnings, and their stop work ages. The overall mean value of SSW_a^* for any given cohort is generally negative at young ages because the present value of taxes to be paid exceeds the present value of benefits to be received, given appropriate discount rates. The mean SSW_a^* becomes positive around age 30, peaks at some age around retirement, and then declines as retirees draw down their accumulated SSW.

The life-cycle patterns of mean expected SSW_a^* are immediately useful for understanding how changes in Social Security policy transfer resources across and within generations, and for understanding overall system finances. Every cohort begins life (at age $a=0$) with a mean value for their expected SSW_0^* , and that average tells us about the extent of intercohort redistribution. At the same time, there are important differences in SSW_0^* within cohorts because benefit replacement rates decline with lifetime income, making the system more progressive. Given the progressive Social Security system parameters, the lower lifetime income members of a cohort begin life with less negative or even positive SSW_0^* , and that is offset by the fact that the average SSW_0^* of the higher lifetime income group in the same cohort is more negative than the overall cohort average. There is some empirical uncertainty about how much SSW_0^* varies within cohorts because the higher replacement rates enjoyed by lower lifetime income participants are offset at least in part by higher mortality rates.⁴

Across and within cohort, net redistribution can be measured at any point in the life cycle, including at age 0, by computing values for mean SSW_0^* within and across cohorts. Most empirical estimates of Social Security progressivity/redistribution use a variant of the SSW concept in the notation above. For example, the Congressional Budget Office redistribution measures are based on the value of taxes paid and benefits received for individuals at retirement, but it is similar to SSW_0^* in the sense that it

⁴ There are also some indications that mortality differentials are rising, which makes the system less progressive. See for example, the National Academies study, *The Growing Gap in Life Expectancy by Income: Implications for Federal Programs and Policy Responses*, available at <https://www.nap.edu/catalog/19015/the-growing-gap-in-life-expectancy-by-income-implications-for>.

captures taxes paid and benefits received over the entire life cycle.⁵ Rather than discount all tax and benefit flows back to age zero, taxes paid and benefits received before age 65 are grown forward by the real discount rate, while taxes paid and benefits received after that are discounted back to age 65. These measures answer the direct question, “as of age 65, how much has an individual paid into the system, and how much will they get back out?” CBO reports the ratio of (discounted) lifetime benefits received to lifetime taxes paid — the so called “money’s worth” ratio — and the two components relative to (discounted) lifetime income. The CBO estimates show that Social Security is indeed progressive, even after controlling for observable mortality differentials. Lifetime benefits received are about twice lifetime taxes paid for the lowest lifetime income quintile, and the ratio of lifetime benefits to taxes falls to about 60% for the highest lifetime income quintile.⁶

One oft-cited example of these sorts of money’s worth or net redistribution measures based on expected SSW is the introduction of the Social Security system itself. The earliest recipients (oldest workers) paid relatively little in tax yet were entitled to the same benefits as the younger workers who would face a lifetime of taxes before they received any benefits. In that sense, those earliest recipients received a substantial wealth transfer from future generations when the system began. In the year the program

⁵ The latest CBO projections are available at <https://www.cbo.gov/publication/55590>. For a more detailed discussion of progressivity in Social Security in published CBO reports, see <https://www.cbo.gov/sites/default/files/109th-congress-2005-2006/reports/12-15-progressivity-ss.pdf>.

⁶ OCACT also publishes annual money’s worth ratios, for hypothetical scaled workers, available at www.ssa.gov/oact/NOTES/ran7/index.html. The OCACT lifetime money’s worth ratios show the same progressive patterns as the CBO estimates.

was put in place, their cohort mean SSW_a^* went from zero to a large positive number, which was in principle offset by large negatives for younger (and unborn) cohorts.⁷

Another overlooked yet more recent example was the substantial expansion of Social Security in the early 1970s, because the birth cohorts who were nearing or already in retirement when the expansion occurred received much higher benefits after the policy change, yet at the same time they had paid less in taxes (relative to the new tax schedules that accompanied the higher benefits) during most of their working years. This substantial intergenerational wealth transfer is likely a direct contributor to the observed patterns of labor force participation of older men in the quarter century following the policy change. Middle-aged and older workers received a substantial positive wealth shock, and they reacted by retiring earlier, likely triggering or exacerbating the decline in labor force participation that lasted through the mid-1990s.

Capturing wealth transfers by looking at expected SSW across and within birth cohorts is also relevant for ongoing policy discussions. There is a direct connection between system solvency and the expected SSW measures described above. When the timespans for computing SSW are limited to the Social Security 75-year valuation period, the population-weighted average expected SSW across in-scope cohort groups is just the overall system actuarial deficit (plus the current trust fund balance). The overall actuarial deficit is currently estimated to be negative. That is consistent with the idea that although workers approaching or already in retirement have large positive

⁷ The offset is “in principle” because the SSW of future generations is an endogenous object that depends on the trajectory for future earnings, population growth, and real discount rates, which may in turn be affected by the policy change (Leimer 2016).

expected SSW, younger and even unborn workers (anyone who will begin working within the 75-year window) have more than offsetting negative expected SSW.

How should we think about expected SSW in a world of projected insolvency? Although there is a great deal of uncertainty about *how* Social Security insolvency will be resolved, there is a legal baseline that can be used as an analytical starting point. The Social Security system has no borrowing authority — the system can only spend what is in the trust fund. Thus, failure to act before the trust fund asset reserves deplete will eventually lead to benefit cuts. In that sense, the expected SSW measures under the no-action (“payable”) scenario should be computed using benefit streams that have those benefit reductions built-in. Although CBO and other observers also report the so-called “scheduled” scenario in their money’s worth calculations, policy analysis of solvent alternatives is appropriately based on comparisons against the payable scenario. In that sense, expected SSW profiles *already* include reductions in benefits after the trust fund asset reserves deplete, to be disproportionately borne by young and even unborn cohorts.

Expected and termination Social Security wealth

Expected SSW captures how much an individual at a given age expects to receive from and pay into the Social Security system going forward and, when aggregated, is useful for characterizing lifetime Social Security redistribution across and within cohorts. There is an alternative concept of SSW suggested by the notation above — the termination value — that is more relevant for thinking about the incentive effects of Social Security over the life cycle (Coile and Gruber 2007; Goda, Shoven, and Slavov 2011). The termination value is the starting point for measuring the net return to working

an additional year, because it is the PDV of benefits earned as of a given age because of work and taxes paid at earlier ages. Measuring the net return (in terms of Social Security) just involves computing the change in PDV of benefits from working an additional year and subtracting the taxes paid in that extra year of work, then dividing the net by current age termination SSW or income.⁸

For our purposes, the key question is which of the two SSW concepts is most appropriate for *completing* existing measures of household wealth inequality? That is, should we be adding expected SSW or termination SSW to the other components of household wealth when constructing a more comprehensive wealth measure for studying inequality? The answer is not obvious, and the underlying reasons why it is not obvious raise more general questions about measuring life-cycle saving, wealth accumulation, and wealth inequality.

One way to think about which concept of SSW is most appropriate for overall wealth inequality is to start with the case where it does not matter, which is the SSW of an individual after they have stopped working. In retirement, expected and termination SSW are identical — both are just the discounted present value of the future benefits the individual will receive until they die. As in a DB pension plan, the individual has earned the right to those benefits, and the individual has no further offsetting obligations in terms of plan contributions or taxes. Wealth is just the discounted present value of those future benefits. In a DB plan, the financial assets held by the plan exactly match those future benefit claims. Assuming future taxpayers will not renege on those

⁸ Indeed, the data set we have developed for this paper is well-suited to empirically investigate the work (and saving) incentives associated with Social Security.

promised benefits, SSW is well-defined, measurable, and equivalent under the two concepts.

Prior to retirement, the values for expected and termination SSW can be different, particularly early in the life cycle, and it could matter which is added to the other balance sheet components to create a comprehensive wealth measure. Again, the appropriate choice is tied to the specific question being asked. For example, one could ask, “as of a given age prior to retirement, how does projected retirement readiness vary within and across generations?” This sort of question is appropriately answered using expected values for SSW (and DB pension benefits) as of retirement age (Jacobs et al. 2019). However, the internally consistent answer to this question also involves expected values for *all* components of household wealth as of retirement age. If the age for evaluating future retirement readiness is (for example) set to 40, the expected retirement readiness analysis involves computing SSW and DB pension wealth based on expected future work through retirement age, because age 40 termination values are uninformative about the wealth that individual will hold going into retirement. The other components of household wealth should also be adjusted using the same expectations about future work, along with some assumptions about saving out of the income from that future work.

An alternative to the retirement-readiness question involves asking the more general question about life-cycle saving and wealth accumulation, “as of a given age, how much wealth has the individual accumulated?” Setting aside SSW for a moment, wealth in this sense is just marketable assets less liabilities, which is exactly the concept built into the household sector balance sheets in the Financial Accounts of the

United States, and exactly the concept used in the literature on household wealth inequality. One component of that household wealth measure is the present value of future DB benefits, and there are marketable financial assets underlying the DB benefit claims.⁹ The marketable assets held by DB plans are consistent with termination values — DB plans are not legally required to hold assets against expected or continuation benefits — and thus the appropriate corresponding discounted present value of DB at the micro level uses the termination concept (Sabelhaus and Volz 2019).

Given the treatment of DB plans, it seems that adding termination values for SSW is also the right answer to the question about how to complete the measures of household wealth for purposes of studying wealth inequality. However, the fact that termination SSW is always zero or positive — an individual will not pay any more taxes if they never work again — highlights the conundrum with using termination values. DB plans have underlying assets to offset the liabilities of the plans to their participants. Yet the wealth claims of individuals with positive SSW have no legal basis without the wealth claims *against* those with negative SSW, and negative SSW is necessarily associated with continued working. Thus, although termination value seems more consistent with studying wealth inequality at a given age, the expected values are more relevant when studying wealth inequality across age groups and time. Thus, most of our focus is on expected SSW in the empirical work, though we do show termination values (mostly for reference and clarification of expected values) at various points.

⁹ Technically, in the Financial Accounts, the value of DB pension claims is the sum of marketable assets held by the plans and unfunded liabilities. The unfunded liabilities are assets of the household sector offset by claims against the sectors (corporate, noncorporate, government) that sponsor the plans.

3. Data and methods

Our goal is to produce empirical estimates of Social Security Wealth (SSW) that can be combined with existing household-level wealth measures and, thus, create more comprehensive wealth measures for studying life-cycle saving, wealth accumulation, and wealth inequality. Achieving this involves starting with high-quality, household-level balance sheet data, then adding the various estimates of SSW as described in the previous section. Our empirical framework makes it possible to create both expected and termination SSW under alternative (scheduled and payable) benefit scenarios and various discounting assumptions. In this section, we describe the microdata, explain how we construct the SSW measures, benchmark those estimates against published Office of the Chief Actuary (OCACT) values, and show how aggregate SSW estimates compare to other household wealth components.

The microdata used here is the Survey of Consumer Finances (SCF) for 1995 through 2019. The SCF is a triennial cross-section focused on household balance sheets. It has extensive information about incomes, demographics, and labor force experiences.¹⁰ The SCF sampling strategy is unique among public-use household surveys, as about one-fourth of the sample is drawn from administrative data records in order to capture the top of the wealth distribution and, thus, overall aggregate wealth.¹¹

¹⁰ See Bhutta et al. (2020) for a discussion of the SCF and the most recent results, for survey year 2019.

¹¹ See Feiveson and Sabelhaus (2019), Bricker et al. (2016), and Batty et al. (2019) for a discussion of how well the SCF captures macro aggregates over time, and how the SCF results on wealth distribution compare to other studies based on a straight read of the administrative data. Another key, but often neglected, advantage to using the SCF is the extensive interconnectedness between income and balance sheet components in the survey instrument that greatly improve data quality. For example, respondents may initially report no

This oversample of high-wealth households, in addition to the detail on household portfolios, makes the SCF uniquely qualified for estimating how adding SSW affects life-cycle saving, wealth accumulation, and wealth inequality. Since the SCF is a series of cross-section snapshots, calculating the PDVs of Social Security taxes and benefits requires estimating life-cycle earnings for individuals and their spouses/partners. Estimating PDVs (for both SSW and DB pensions) requires individual mortality rates, which we differentiate by age, sex, income, and birth year.

Life-cycle work and earnings

The first crucial data development step is to use SCF detailed work history and work expectations modules to create life-cycle earnings measures from the cross-section observations, because life-cycle earnings are the key to PDVs of taxes and benefits. The SCF data on labor force participation, employment, unemployment earnings, current job characteristics, past jobs, and expected future work are collected in detailed modules for both the respondent and spouse/partner (if present).¹² The sequence of questions in the SCF labor modules focus first on current employment

owned businesses, but then later in the survey reveal that they think of themselves as self-employed, which triggers reconsideration of the owned business questions. Similarly, the existence of various types of DB pensions and DC/IRA accounts is checked in various ways throughout the survey, as respondents are asked at various points about coverage, balances, and benefits/withdrawals in different ways. The extensive SCF case review process undertaken by Federal Reserve Board staff is largely focused on sorting through what is sometimes conflicting information about these complicated balance sheet items in the raw data.

¹² In SCF microdata files the respondent is always recoded to be the man in a different sex couple, and the individual closest to 40 years old in a same-sex couple. This underscores the use of data for both respondent and spouse/partner in the empirical work here, because there are many cases where the spouse/partner will be the primary earner.

status, then (if employed) current main job characteristics, including pensions. That is followed by modules on employment and job histories and expected future work.

The idea of working versus not working is a simple idea for most survey respondents, but there are a variety of circumstances that can complicate the issue for some. Many respondents have multiple statuses: They are both working (full or part time) and a student, homemaker, volunteer, or some other status that is generally considered out of the labor force. The SCF instrument lists the possible situations, and then collapses (for the purpose of determining survey question sequence) people into three bins: not working, working full time, and working part time. The information one would like about the individual varies with the bin, and the SCF survey instrument is tailored to the three groups accordingly.

If the individual reports working full or part time, they are asked a number of questions about their current main job.¹³ Many of these questions are standard in household surveys, including whether the respondent is self-employed or working for someone else, how much they earn through a regular salary and/or variable pay, their usual hours and weeks worked, their industry and occupation, and the size of the firm they work for. The key incremental information about their current job, relative to most cross-sectional data sets, that makes it possible to start building the longitudinal variables is that respondents are also asked when they started the current job and when they expect to stop working on the current job.

¹³ The survey also collects limited information about second jobs, but that is not used for purposes of classifying work status.

After collecting the current job data, the survey turns to work histories. There are three distinct sequences of questions for individuals who report not working, working full time, and working part time, respectively. The survey instrument's goal is to retrieve some common work history meta data for the three groups, including the total number of years worked full time and part time, and details about the longest full-time job ever worked. The "longest job" questions are a subset of the information collected about the current main job, including salary in the last year they worked the job, industry, occupation, and self-employed versus working for someone else. The question sequences vary because the point of reference differs by current work status. For example, the leading question "have you ever worked full-time?" will only make sense to respondents who are currently not working or working part-time.

Even with carefully tailored question sequences, there is still some respondent confusion about and disagreement among the various work history responses that must be addressed in the data construction. For example, one common tendency is for respondents to confuse current main job "type" with their current employer. If they have been working the same type of job their whole adult life, they will sometimes answer that they started the current job at a young age, even though the earlier job(s) they are thinking about involved a different employer. The self-employed are particularly prone to answering they have been in the current job since childhood, though in their case some of them have indeed worked in family businesses all their lives. Building the longitudinal work histories requires a balancing of sometimes conflicting information from the current

job, the longest full-time job, and the overall number of years worked full and part time.¹⁴

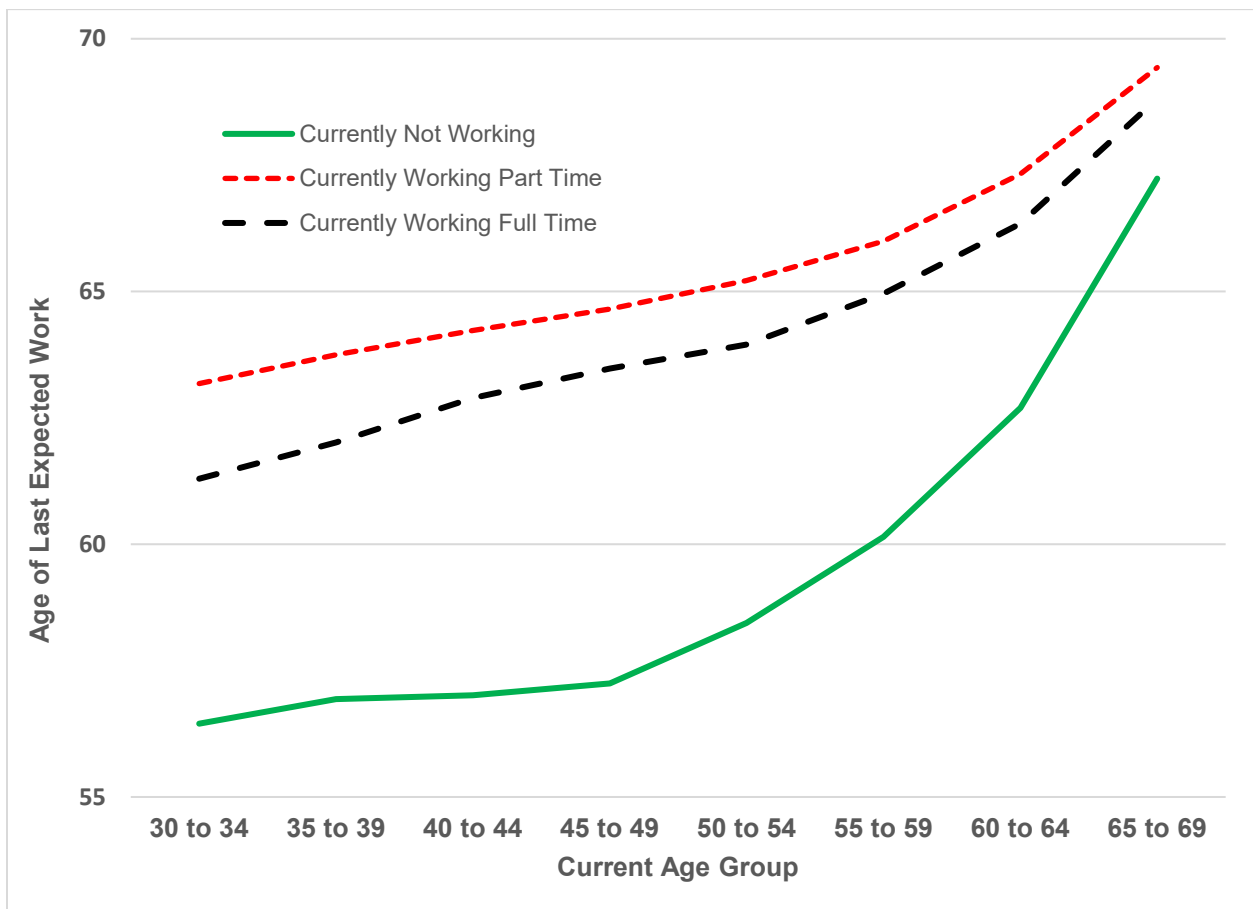
Forward-looking measures are based on what respondents report they expect to do in the future and, thus, the questions on future work are crucial. Again, the specific questions vary with current work status, but the general idea is to divide the future into periods of expected full time, part time, and not working. The majority of the working population — those who are working full time in the survey year — are the easiest, because the survey instrument simply focuses on when they expect to stop full-time work, whether they expect to switch to part-time after that and, if so, when they expect to stop working completely. The currently not working and working part time are first asked if they expect to start working full time before the rest of the question sequence kicks in. Again, there are sometimes conflicting answers *vis-a-vis* the current main job, situations where the respondent initially says they “never expect to stop” doing the particular type of job they are in, but then later say they actually do expect to stop at a particular age.

Reconciling those conflicting answers is key for generating expected stop work wages. Using the reconciled data, expected retirement of well-attached workers lines up well against observations from other data sets and the actual (recalled) retirement behavior of older respondents in the SCF itself (Sabelhaus 2019). The specific concept of last-work age we use in the expected SSW calculations is somewhat dependent on the respondent’s current work status because retirement and Social Security benefit

¹⁴ The code that reconciles the various employment history variables is available from the authors upon request.

claiming are imperfectly correlated. If the respondent currently works full time or expects to become full time in the future, their expected last-work age is when they say they will stop working full time, because many respondents state they expect to transition to part-time work at the end of their career. If they are currently part-time, then expected last-work age is the age at which they expect to completely stop working. All answers are capped at 70, the age Social Security benefits no longer accrue.

Figure 1. Respondent-reported expected last-work age



Notes: Authors' calculations using 1995 through 2019 SCF. Conditional on not yet receiving Social Security benefits.

The reconciled expected last-work ages vary predictably across current age and current work status (Figure 1). Younger workers tend to give lower ages for expected stopping work than older workers with any current work status. The average expected stop work age increases with age, because expectations become more realistic, and because some transition from working to not working or claiming benefits (those already claiming are not in Figure 1). Also, expected last-work age at any given age (younger than 65 to 69) for full-time workers increased, on average, about two years between the 1995 and 2016 survey waves, which is in line with actual and projected Social Security benefit claiming patterns.¹⁵ Thus, our expected SSW calculations capture the phenomenon that younger cohorts at a given age and (real) income will have higher benefit PDVs (given the same discount rate) because they expect to work longer. Whether or not the higher benefit PDVs translates into higher expected SSW depends on taxes and mortality.

Calculating Social Security benefits requires earnings at each point in the life cycle. Earnings are assigned at each age for which an individual is predicted to be working either full time or part time. We estimate an earnings equation that is a function of age and its square to provide age-specific for each sex/education combination. This provides a sex- and education-specific growth rate for each age from 18 through 70. For individuals working at the time of interview, their current wage is used to project future wages through expected retirement, using the growth rates described. The wage is further adjusted based on their current and projected future work status using a simple scalar, i.e., expected full-time to part-time transitions involve a 50% earnings reduction.

¹⁵ See, for example, Dudel and Myrskylä (2017).

Retrospective earnings are based on the “longest past job” questions in the SCF. Workers are asked about how much they were earning (and at what age) when they left their most important career job. We use that information along with current earnings to pin down a life-cycle earnings trajectory. We use the same age-specific growth rates as above to trace out expected earnings over the tenure of one’s longest past job. If an individual does not have a long past job, his survey earnings are also used to predict earnings earlier in the life cycle. For an individual with a long past job, the final earnings from the long job are used to fill in ages preceding the beginning of the long job. For ages between the end of a long job and the beginning of the current job, a linear mapping is used to connect the long job final salary and the predicted start salary on a current job. If no earnings are reported in the survey, which is especially problematic for those who report never having worked full time, we assign median earnings for ages 20 to 24 in survey year as a baseline if an individual is younger than 30, or the minimum wage for those at other ages (who expect to work in the future). As in the prospective earnings calculations, life-cycle earnings trajectories are adjusted for changes in work status.¹⁶

¹⁶ Benchmarking our estimated PDVs of taxes and benefits against OCACT aggregates (see Appendix) suggests that our assumptions about earnings may be conservative on net, though it is not clear how much we are missing because of limiting our calculations to SCF respondents and spouses. The SCF essentially captures aggregate taxable earnings in the survey year, but the PDVs are low, suggesting our assigned earnings are too low at other points in the life cycle. This is an important area for further research.

Demographics and differential mortality

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. 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.¹⁷ 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. Appendix B explains how the Chetty et al. (2016) mortality differentials are applied to the SCF data in a way that preserves average mortality from the Social Security Administration and *relative* mortality by income across age, sex, and cohort population groups.¹⁸

¹⁷ The SSA mortality data is available at www.ssa.gov/OACT/HistEst/Death/2019/DeathProbabilities2019.html.

¹⁸ Differential mortality is explicitly incorporated in the Social Security actuarial projections as well, with an adjustment based on career average earnings. See www.ssa.gov/oact/NOTES/pdf_studies/study124.pdf.


Actual versus predicted benefits

Estimated life-cycle earnings and expected last-work ages are the main inputs into the SSW calculator. We compute own worker benefits based of the expected retirement last-work age, then initiate the benefit stream at $\max(62, \text{stop work age}+1)$. We also calculate expected spouse/survivor benefits for married couples, and the PDVs are computed using own and spouse/partner differential mortality.

The approach to calculating SSW raises the question about choosing our computed benefits versus actual respondent reported values when the actual benefits are available. Across survey waves, the fraction of respondents for whom we have actual benefit values (and benefit start ages) increases predictably with age (Figure 2). In the calculations below, we use the actual benefit values, including those for retirement, spousal, survivor, and disability when they are available. Reported benefit aggregates in the SCF for respondents and spouse/partners were about 86% of actual benefits paid in 2018.¹⁹ We consider that ratio quite good, given that we are missing benefits paid to individuals outside the SCF sample frame (especially those living in institutional settings) or residing in an SCF household where they are not the respondent or spouse/partner. Still, it is possible that some respondents are making systematic errors, such as reporting benefits after Medicare premiums are deducted.

¹⁹ Another important exercise for future work that will likely help refine our life-cycle earnings estimates is to compare the actual and predicted earnings for respondents who are already claiming.

Figure 2. Fraction already receiving benefits by current work status



Notes: Authors' calculations using 1995 through 2019 SCF.

DB pension wealth

The published SCF wealth concept — based on balance sheet components directly measured in the survey — does not include SSW or DB pension wealth. As with Social Security, the wealth equivalent of a stream of future DB pension incomes is a survival-adjusted present value. Those wealth equivalents are constructed using survey information about currently received pension benefits, expected future pensions, and pensions associated with current jobs, using the same differential mortality model and discounting assumptions applied to SSW above. The sum of all pension wealth is

benchmarked to the aggregate from the Financial Accounts of the United States (FA), and pension wealth distribution is benchmarked against direct measures from the HRS. Details of the DB wealth estimates used here follow (see also Sabelhaus and Volz 2019).

The SCF collects comprehensive, detailed information about retirement plans associated with current and past jobs of the respondent and their spouse or partner. The measure of wealth consistent with the comprehensive household balance sheet is the present value of those 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 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.

The survey questions about currently received benefits and expected benefits from past job pensions are standard. The respondent is simply 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 value calculation that also involves an assumed interest rate and respondent expected (and demographically differentiated) mortality. This present value calculation is the level of financial assets that the retirement plan sponsor must hold to pay promised benefits and, thus, corresponds directly to the household wealth measure we are trying to

capture. In total, these present value calculations for SCF respondents indicate that about 40% of the aggregate DB pension assets in the FA is attributable to the promised benefits of currently receiving and past job pensioners, and thus the remainder is accumulated for workers still on the jobs for which they are accumulating the rights to future DB benefits.

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, and, in any event, that eventual benefit is likely different from the FA benchmark termination concept. The SCF asks a series of questions about the current job's DB pensions, including how long the worker has been in the plan, when they expect to receive benefits, and how much they expect to receive after benefit payments begin. As with SSW, these questions make it possible to construct a few different measures of what the stream of future benefits represents in present value. For example, it is possible to compute the present value of the future benefit stream assuming the worker remains in the job until his or her expected retirement age, for a given projected final salary, and under the assumption that the worker knows and reports values consistent with the actual benefit formula in the survey. DB benefit formulas are generally based on a complicated combination of years of service and average "high" salary, and those formulas are generally beyond the grasp of survey participants. Of course, if constructed properly, this "continuation" value of the DB pension is an important and useful concept when thinking about life-cycle financial planning.

The continuation value represents what the DB pension *could* be worth to the worker at some point, based on a series of assumptions about future employment and earnings under the current plan rules. However, the measure of DB wealth that corresponds to total household wealth in the FA is the narrower financial liability of retirement plan sponsors: The DB wealth that the worker has accumulated to date. Plan sponsors are not required to set aside the continuation value of a pension for every worker in their plans. Rather, DB plans are required to hold only the present value of benefits already earned by the worker, the “termination” value, which corresponds to DB wealth in the FA. 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 termination value of a DB pension is always less than or equal to the continuation value, and the two converge as the worker approaches retirement age, at which point they are equal.

Our approach to solving for DB pension termination values for current workers in the SCF relies on three complementary sources of information. The first piece of information is the aggregate value of DB pension assets from the FA. After subtracting the roughly 43% of total assets accounted for by currently received and known future benefits (as described above), the remainder represents the legal claims (termination values) of current-job DB participants. The second set of inputs, from the SCF, is the worker’s age, the number of accumulated years of plan coverage, earnings, and sector (private or public) of the job held by the worker. Finally, the calculation involves data from Fang, Brown, and Weir (2016) from HRS to validate our estimated person-level termination values. After building in a wedge between public and private sector DBs

(public plans are more generous than private, relative to earnings), the actuarial present value calculations used to construct termination values are based on the same differential mortality model we also use for estimating SSW.

Aggregate household wealth

In 2019, SCF published household net worth was roughly \$96 trillion, up from \$21 trillion in 1995. As a prelude to the distributional and life-cycle analysis of the expanded wealth measures that include DB and SSW, it is useful to put the three aggregates side by side over time. How large are DB wealth and SSW relative to published SCF net worth? Are those ratios changing over time? How is the ratio of each evolving relative to aggregate income? In doing so, we establish important facts about the various constructs that are key for the distributional analysis later. For example, although aggregate SSW is large relative to other wealth components in every year of our sample, it is not growing as fast as SCF published net worth. The size and growth of SSW depends on whether one looks at expected or termination SSW, the choice of discount rates, and whether one considers scheduled or payable benefits.

In 2019, aggregate net worth in the SCF was about \$96 trillion, and aggregate DB pension wealth was about \$19 trillion (Table 1). The estimates of aggregate SSW in 2019 ranged from as low \$19 trillion to as high as \$40 trillion, depending on the specific wealth concept — expected or termination — along with the assumptions about what rates to use for discounting future taxes and benefits, and whether future benefits are as scheduled or limited to payable under current OCACT projections. Termination SSW is always above expected SSW. Payable benefits are always lower and declining as a share of scheduled benefits over time, because we are getting closer to insolvency as

we move through the sample period. Finally, choosing OCACT real discount rates instead of the fixed 2.8% real rate is generally benign before 2016.

Most of our focus in what follows is on expected SSW, but it is worth noting here that aggregate termination SSW is notably higher than aggregate expected SSW in all years. Under the expected SSW concept, on average, the PDV of benefits reflects the higher benefit the worker will receive because of additional years worked, but that comes at the cost of additional taxes. Careful consideration of the Social Security benefit formula makes it clear why termination SSW is higher for many participants, especially the young. A young worker with high lifetime earnings will have a low average indexed monthly earnings (AIME) because they have many zeroes in their work history. If they stop working when young, they will be entitled to the benefit replacement rate of a lifetime low earner, which is a higher rate of return. Additional years of work raises their AIME and reduces the rate of return on additional taxes paid.

The substantial difference between payable and scheduled expected SSW is growing over our sample period, but not because of changes in assumptions about Social Security Trust Fund asset reserves depletion. The payable scenarios for every survey wave assume that 80% of benefits will be payable after 2034, consistent with recent *Trustees Reports*. The ratio of payable to scheduled expected SSW fell from about 90% to 80% after 1995. Barring policy change, the ratio will continue falling as we get closer to 2035. Indeed, it is worth noting that 80% is not the lower bound on payable relative to scheduled SSW: 80% is the long-run ratio of payable benefits relative to scheduled benefits. The payable scenarios assume that taxes are still collected as

scheduled, and with the PDV of taxes at roughly one-third the PDV of scheduled benefits, the long-run projected ratio of scheduled to payable SSW is more like 50%.

The rate used to discount future taxes and benefits turns out to be important for SSW, especially in this current era of low real interest rates. In general, our fixed 2.8% real rate specification lines up well with the year-by-year OCACT real discount rates. This is to be expected because 2.8% is the average real discount rate used by OCACT over this period. The similarity breaks down after 2013, however. Our estimated SSW grows systematically (relative to aggregate income) for life-cycle reasons in the second half of the sample period as the Baby Boom generation ages and approaches benefit claiming, but the jump in SSW after 2013 using the OCACT discount rates greatly magnifies that increase. After 2013, OCACT lowered their assumed long-term discount rate, but most of the effect on PDVs comes from changes in the OCACT assumptions about how quickly interest rates transition from low current values to the assumed long-run rate.

The discount rate's impact on the PDVs is substantial. Indeed, if taken to logical extremes, it effectively drives the relationship between SSW and other household wealth. That, in turn, drives how including SSW affects overall wealth inequality. If one pegs the discount rate to move in step with, say, observed market interest rates, it is possible to generate very rapid growth in aggregate SSW (Catherine, Miller, and Sarin 2020). In our framework, similar aggregate SSW growth for our sample period can be generated using a 1% real discount rate after 2013, which leads over \$40 trillion in SSW in 2019, nearly double the baseline estimate. Because adding SSW is always wealth-

equalizing in levels, that much additional growth in SSW is basically enough to offset rising concentration in the other components of household wealth.²⁰

The choice of discount rate is a philosophical conundrum because there is obviously no market in which future Social Security tax liabilities and benefit claims are traded. SSW is clearly unlike the financial instruments for which interest rates have fallen in recent years, as the government can change taxes and benefits in response to system financing problems. Acknowledging that unique (and increasingly relevant) risk would likely offset the declining time value of money in a complete discounting framework. The real question is how the choice of discount rate affects life-cycle behavior, which is beyond the scope of this paper.

²⁰ A direct comparison with Catherine, Miller, and Sarin (2020) is not possible because the empirical framework for combining SCF wealth and estimated SSW is fundamentally different. We use the SCF microdata to generate earnings histories and expectations, which are used to estimate SSW, and we use our estimates of DB wealth from Sabelhaus and Volz (2019). Catherine, Miller, and Sarin (2020) use another microdata set to estimate aggregate SSW, then splice the estimates into the published SCF wealth distributions.

Table 1. Wealth aggregates by year (billions)

	1995	1998	2001	2004	2007	2010	2013	2016	2019
SCF Published Net Worth	\$ 21,101	\$ 29,087	\$ 42,338	\$ 50,354	\$ 64,649	\$ 58,207	\$ 64,748	\$ 86,865	\$ 96,073
+ Defined Benefit Wealth	\$ 5,943	\$ 7,044	\$ 8,457	\$ 10,031	\$ 11,923	\$ 13,474	\$ 15,904	\$ 17,683	\$ 19,092
= Household Wealth	\$ 27,044	\$ 36,131	\$ 50,795	\$ 60,385	\$ 76,572	\$ 71,681	\$ 80,651	\$104,548	\$115,165
Baseline Scheduled Benefits Scenario									
Expected Social Security Wealth	\$ 6,469	\$ 7,573	\$ 9,507	\$ 11,529	\$ 14,177	\$ 16,180	\$ 19,382	\$ 21,589	\$ 23,931
= PDV Expected Benefits	\$ 12,494	\$ 14,525	\$ 17,614	\$ 20,549	\$ 24,407	\$ 27,168	\$ 31,139	\$ 34,246	\$ 38,254
- PDV Expected Taxes	\$ 6,025	\$ 6,952	\$ 8,107	\$ 9,021	\$ 10,230	\$ 10,988	\$ 11,757	\$ 12,657	\$ 14,323
Termination Social Security Wealth	\$ 9,649	\$ 11,246	\$ 13,782	\$ 16,361	\$ 19,455	\$ 21,929	\$ 25,497	\$ 28,269	\$ 31,378
Baseline Payable Benefits Scenario									
Expected Social Security Wealth	\$ 5,904	\$ 6,758	\$ 8,354	\$ 9,968	\$ 12,054	\$ 13,517	\$ 16,011	\$ 17,569	\$ 18,895
= PDV Expected Benefits	\$ 11,929	\$ 13,709	\$ 16,461	\$ 18,989	\$ 22,285	\$ 24,505	\$ 27,768	\$ 30,226	\$ 33,217
- PDV Expected Taxes	\$ 6,025	\$ 6,952	\$ 8,107	\$ 9,021	\$ 10,230	\$ 10,988	\$ 11,757	\$ 12,657	\$ 14,323
Termination Social Security Wealth	\$ 9,423	\$ 10,907	\$ 13,244	\$ 15,581	\$ 18,330	\$ 20,439	\$ 23,471	\$ 25,708	\$ 28,018
Scheduled Benefits Using OCACT Discount Rates									
Expected Social Security Wealth	\$ 6,656	\$ 6,541	\$ 8,489	\$ 10,327	\$ 13,230	\$ 14,891	\$ 18,561	\$ 27,575	\$ 32,846
= PDV Expected Benefits	\$ 12,361	\$ 13,104	\$ 16,333	\$ 19,012	\$ 23,238	\$ 25,484	\$ 30,273	\$ 42,710	\$ 50,484
- PDV Expected Taxes	\$ 5,705	\$ 6,564	\$ 7,844	\$ 8,685	\$ 10,008	\$ 10,593	\$ 11,712	\$ 15,135	\$ 17,637
Termination Social Security Wealth	\$ 9,402	\$ 10,281	\$ 12,973	\$ 15,358	\$ 18,703	\$ 20,785	\$ 25,035	\$ 34,813	\$ 40,432
Addendum: Permanent Income	\$ 4,461	\$ 5,258	\$ 6,916	\$ 8,020	\$ 9,401	\$ 9,843	\$ 10,496	\$ 12,673	\$ 13,253

Notes: Author's calculations using 1995 through 2019 Survey of Consumer Finances. Baseline scenarios set fixed real discount rate to 2.8% for all years. OCACT discount rates are from the annual data underlying each year's *Trustees Report*.

4. Impact of SSW on top wealth shares

A primary motivation for this work is understanding how including SSW affects estimated levels and trends in wealth inequality. Any distributional analysis of that sort requires that we make decisions about how to sort and weight observations. There are three key decisions, (1) whether to sort the entire population or sort within age groups, (2) whether to weight observations to keep a given number of households versus a given number of persons in a distributional group, and (3) whether to re-sort when the wealth measure is more expansive. Our sorting and weighting approach intends to maintain consistency between overall and within-age inequality, and to purge the effects of evolving living arrangements (both over the life cycle and over calendar years). Thus, we sort within an individual's age group and weight by persons (respondents and spouse/partners), both of which decrease levels of a given top wealth share but do not affect trends. The decision about whether to re-sort after adding SSW turns out to have little impact on estimated top wealth shares.

These decisions about sorting and weighting have a substantial impact on estimated SCF wealth shares in every survey year, but the trend is largely unaffected (Table 2). The reference point for what follows is the published SCF top 1% and top 10% wealth shares, shown in the first line of the top and bottom panels of Table 2, respectively. The published SCF wealth shares are based on sorting households (not persons) across the entire population (not within age groups). Under those sorting decisions, and using the narrow published SCF wealth concept, the top 1% wealth share increased about 4 percentage points over our study period, and the top 10% wealth share increased almost 10 percentage points.

Adding defined benefit (DB) pension wealth to the published SCF top shares lowers the estimated *levels* of wealth concentration but does not reverse the *trend* toward greater wealth inequality (second line in each panel of Table 2). Using the same household-weighted sorting approach with the household wealth measure that includes DB pension wealth, the top 1% wealth share falls from about 35% to about 28% in 1995, but still rises substantially over the study period. Likewise, the top 10% wealth share falls from about 68% to about 61% in 1995, but again, increases 10 percentage points. Indeed, these can be interpreted as deepening concerns about wealth inequality, because the same percentage point increases are being applied to a lower base, so the proportional increases are larger.

Under the traditional household weighted sorting irrespective of age, the share of SCF household wealth (published plus DB) held by the top 10% increased from about 61 to about 71% between 1995 and 2019. When we sort households within their own (10-year) age groups — effectively capturing the top 10% of each age group in each year — the top 10 wealth share increases from about 56% to about 66%. When, in addition, we change the weighting to capture a fixed percentage of respondents and spouse/partners in each wealth group (irrespective of living arrangements), the top 10 share rises from about 53% in 1995 to about 62% by 2019. The two sorting and weighting adjustments have the expected effects at each point in time, as the young rich are less rich than the old rich, and the rich are more likely to be married. Although the decisions about how to sort and weight are contentious from a normative perspective, our decision to go with the more leveling assumptions is driven by our goal of

comparing wealth across and within age groups over time using a consistent frame. The various measures show similar trend increases.

For completeness (and reader choice), we show the top wealth shares including SSW under all three weighting and sorting permutations, and in addition, we show the effects of re-sorting by the expanded total wealth measure. The impact of adding SSW is substantial, dramatically reducing estimated top wealth shares in every year, as expected. However, the percentage point increases in top wealth shares are similar and, thus, as before, the *relative* increase is even larger, because the base is lower. Again, our preferred weighting and sorting approach is the “person-weighted within age groups” because it eliminates the confounding effects of life-cycle wealth (we are not grouping the young low wealth with the old low wealth; they are low wealth for different reasons) and the effects of marital trends (marriage rates have fallen more for the less wealthy).

Under our preferred sorting and weighting approach, the estimated top 10% comprehensive wealth (household wealth plus SSW) share increased from about 45% in 1995 to about 54% in 2019, a 20% increase. The final decision is whether to re-sort by the comprehensive wealth measure. The good news from our perspective is that it does not matter — the estimated wealth shares are basically unaffected. Therefore, we choose to sort by household wealth for the distributional analysis that follows. The fact that re-sorting does not matter simply shows that — after controlling for age and family size — adding SSW does not meaningfully change individuals’ relative positions in the wealth distribution.

Table 2. Estimated top wealth shares

	Top 1 Percent Wealth Shares								
	1995	1998	2001	2004	2007	2010	2013	2016	2019
SCF Published Wealth Shares									
Household Sorting (by Published Wealth)	34.8%	33.8%	32.1%	33.2%	33.6%	34.1%	35.5%	38.5%	37.2%
Household Wealth Shares									
Household Sorting	27.7%	27.9%	27.4%	28.1%	29.0%	28.8%	29.0%	32.5%	31.5%
Household Sorting within Age Groups	24.7%	25.1%	24.0%	25.0%	26.0%	25.1%	25.9%	28.9%	28.1%
Person-Weighted Sorting within Age Groups	23.7%	23.4%	22.5%	23.4%	24.0%	22.7%	23.9%	26.8%	25.9%
Household + Social Security Wealth Shares									
Household Sorting	22.9%	23.6%	23.5%	24.1%	24.9%	24.0%	23.9%	27.4%	26.6%
Household Sorting within Age Groups	20.3%	21.1%	20.5%	21.3%	22.3%	20.8%	21.2%	24.3%	23.6%
Person-Weighted Sorting within Age Groups	19.4%	19.6%	19.2%	19.9%	20.5%	18.8%	19.6%	22.5%	21.8%
Person-Weighted within Age Groups, Resorted	19.3%	19.6%	19.1%	19.9%	20.5%	18.9%	19.5%	22.5%	21.8%
	Top 10 Percent Wealth Shares								
	1995	1998	2001	2004	2007	2010	2013	2016	2019
SCF Published Wealth Shares									
Household Sorting (by Published Wealth)	67.9%	68.6%	69.6%	69.4%	71.4%	74.4%	75.0%	77.1%	76.5%
Household Wealth Shares									
Household Sorting	61.5%	63.3%	64.1%	64.5%	67.1%	68.5%	68.8%	71.5%	71.3%
Household Sorting within Age Groups	55.5%	58.3%	59.5%	59.0%	62.7%	62.2%	63.1%	66.0%	65.9%
Person-Weighted Sorting within Age Groups	52.7%	55.6%	56.6%	56.1%	59.8%	59.2%	59.6%	62.8%	62.3%
Household + Social Security Wealth Shares									
Household Sorting	54.7%	56.7%	57.9%	58.2%	60.3%	60.3%	60.2%	63.6%	63.4%
Household Sorting within Age Groups	48.0%	51.0%	52.9%	52.1%	55.4%	53.8%	54.1%	57.6%	57.6%
Person-Weighted Sorting within Age Groups	45.3%	48.4%	50.0%	49.4%	52.6%	50.9%	50.9%	54.5%	54.1%
Person-Weighted within Age Groups, Resorted	45.1%	48.4%	50.0%	49.2%	52.5%	50.8%	50.7%	54.4%	54.2%


Note: Authors' calculations using Survey of Consumer Finances (SCF). Household wealth is SCF published wealth plus Defined Benefit (DB) pension wealth. Social Security Wealth (SSW) estimates constructed using baseline 2.8 percent real discount rate. Unless otherwise noted, wealth concept for sorting is household wealth.

5. Wealth inequality across and within age groups

The first part of the distributional analysis using our expanded wealth measures focuses on wealth inequality across and within age groups. We compare real average wealth holdings by age in 1995 and 2019, first for entire age groups, then for the bottom 50% wealth group, the 50th through 90th percentile wealth group, and the top 10% wealth group within each age group. As described in the previous section, the wealth groups are person-weighted and sorted within age groups, using the household wealth (SCF published plus DB) wealth concept. We focus on three main wealth components: nonretirement wealth, retirement wealth (DC and DB), and SSW. We show how each component of wealth contributes to the overall age group average, and the extent to which the contribution of expected SSW to overall wealth varies between the scheduled benefits and payable scenarios.


We begin the decomposition by looking across age groups in 1995 and 2019 (Figure 3). The columns stack average wealth by type — expected SSW, retirement, and nonretirement wealth — in 1995 and 2019. The third bar in each age group is the 2019 value with expected SSW reduced according to the payable scenario. In general, the main takeaways from the overall age group averages in Figure 3 are (1) average nonretirement wealth is the largest component, but the relative importance varies with age; (2) overall average wealth increased across all age groups between 1995 and 2019; (3) average wealth increased much more at older ages; and (4) the impact of assuming payable benefits is only very noticeable (against the backdrop of a comprehensive wealth measure) at younger ages.

Figure 3. Mean wealth by age, 1995 and 2019



Notes: Author's calculations based on 1995 and 2019 SCFs. Retirement wealth includes DB and DC plans. Social Security Wealth (SSW) is the baseline expected concept, using real discount rate of 2.8%.

Figure 4. Mean wealth by age, 1995 and 2019, bottom 50% wealth group

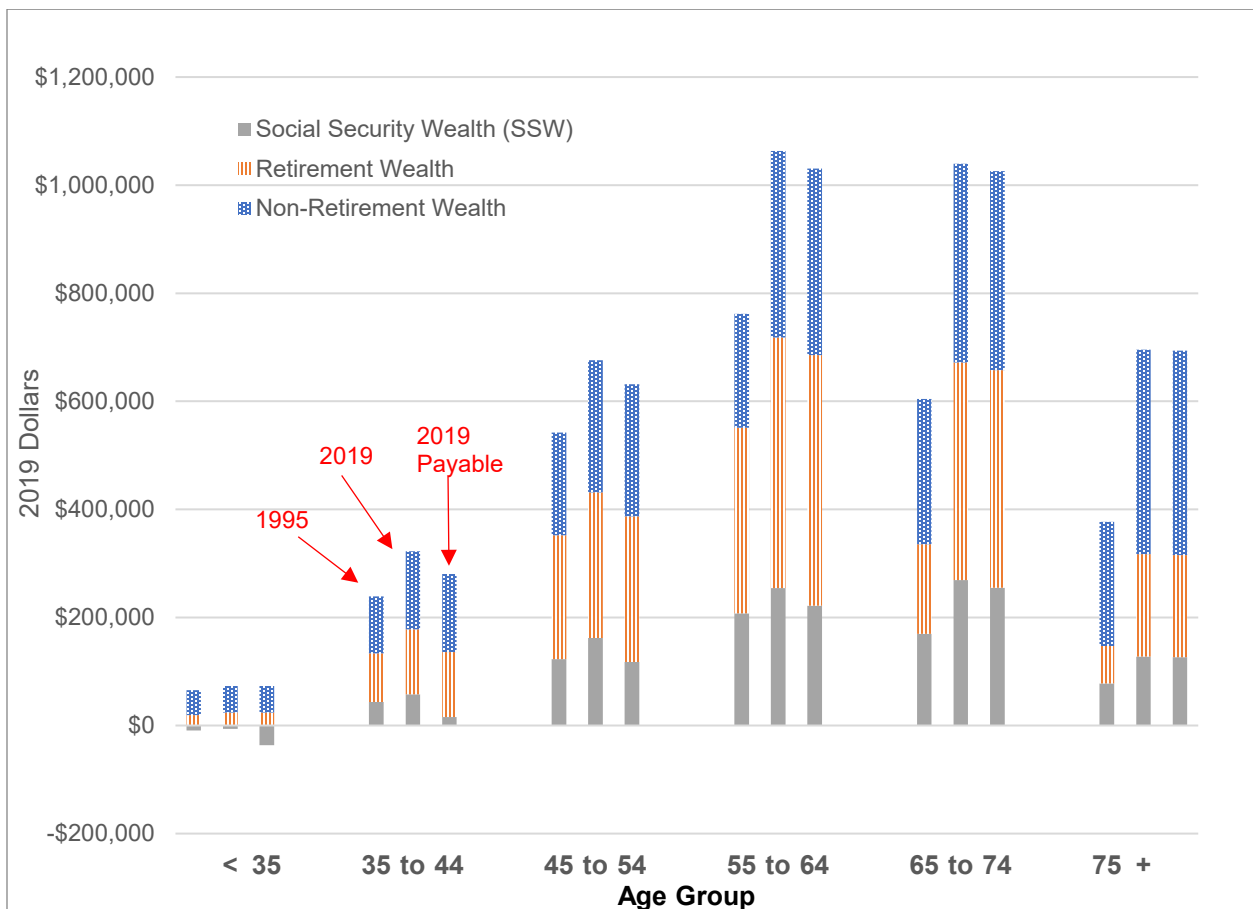


Notes: Author's calculations based on 1995 and 2019 SCFs. Retirement wealth includes DB and DC plans. Social Security Wealth (SSW) is the baseline expected concept, using real discount rate of 2.8%.

Our within- and across-age group wealth inequality analysis is motivated by a growing realization that the gap between older (wealthier) and younger (less wealthy) wealth groups may be growing much faster than the gap between high and low wealth overall or within a given age group. This is borne out when looking within age groups by wealth (Figures 4, 5, and 6). Younger people in the bottom 50% wealth group have seen declines in average total wealth over the past quarter century (Figure 4). The

bottom 50% at older ages have seen wealth gains, though much of that is because of increased expected SSW. Indeed, average wealth for the bottom 50% at most ages is dominated by expected SSW, and the payable SSW bars, show, for example, that switching from scheduled to payable SSW in the 55 to 64 age group eliminates almost all of the gains in expected SSW over the 25 year period, reinforcing the net decline in average wealth associated with the other two components.

Figure 5. Mean wealth by age, 1995 and 2019, 50th to 90th percentile wealth group

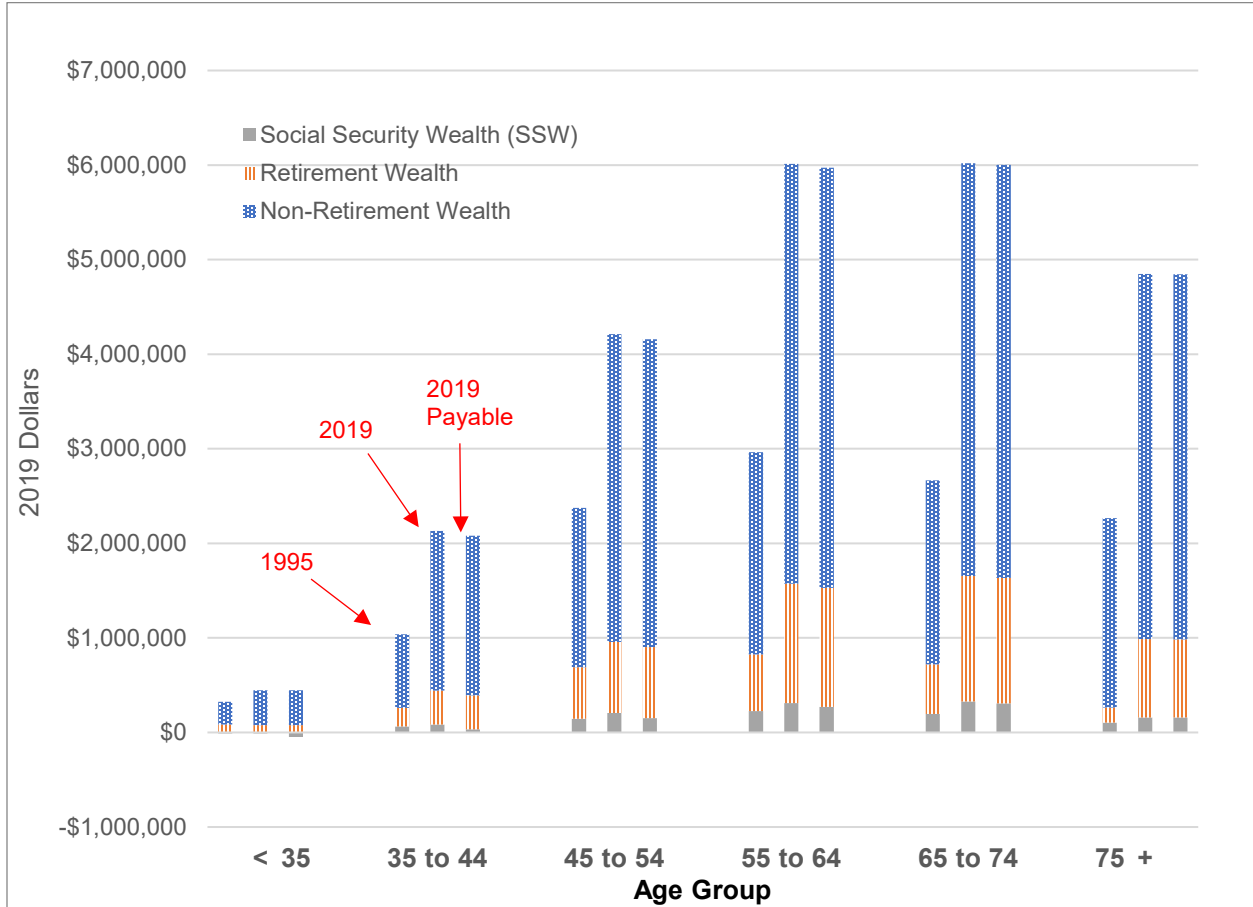


Notes: Author's calculations based on 1995 and 2019 SCFs. Retirement wealth includes DB and DC plans. Social Security Wealth (SSW) is the baseline expected concept, using real discount rate of 2.8%.

Average wealth in the 50th through 90th percentile wealth group has risen much more at older ages, though again only modestly so or even falling at the youngest ages (Figure 5). It is noteworthy that again, increases in expected SSW — mostly driven by higher lifetime earnings leading to higher benefits — are increasing average wealth in many age groups. However, acknowledging that those benefits are not expected to be payable eliminates most of that gain for younger cohorts. It is also interesting to note that retirement wealth (the sum of DB and DC wealth) accounts for a substantial share of wealth for those in the 50th to 90th percentile wealth group, especially those approaching retirement. Nonretirement wealth is also increasingly important for the 50th to 90th percentile wealth group at older ages, likely being driven by rising house values.

The increase in average wealth is proportionally greatest for the oldest age groups in the top 10% of their respective wealth distributions (Figure 6). Expected SSW is barely noticeable on the top 10% column charts, and retirement wealth is also proportionally smaller. There are substantial percentage increases in average wealth for the top 10% by wealth in all age groups, clearly being driven by nonretirement wealth in the form of financial assets, closely held businesses, and real estate. The percentage gains are largest for the top 10% of wealth within the older groups, further reinforcing the old/wealthy versus young/less wealthy inequality narrative.

Figure 6. Mean wealth by age, 1995 and 2019, top 10% wealth group



Notes: Author's calculations based on 1995 and 2019 SCFs. Retirement wealth includes DB and DC plans. Social Security wealth (SSW) is the baseline expected concept, using real discount rate of 2.8%.

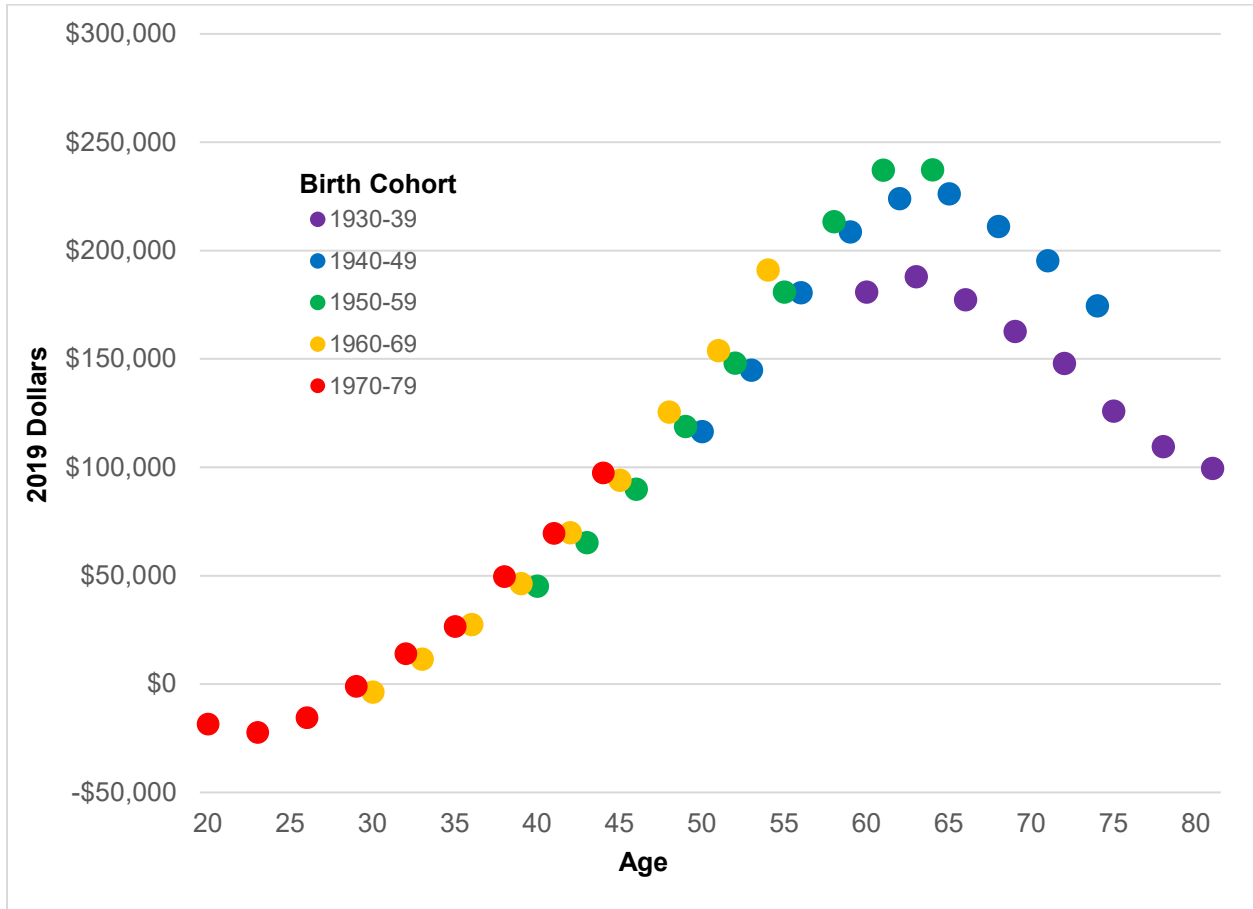
6. Life-cycle patterns of Social Security wealth

The advantage of constructing longitudinal earnings for multiple SCF waves is the ability to follow the same birth cohorts over time in a pseudo-panel framework, which, in turn, makes it possible to construct life-cycle wealth profiles. The triennial structure of the SCF gives us up to nine pseudo-panel SSW observations per 10-year birth cohort over the quarter century span between 1995 and 2019. By connecting the cohort averages between survey waves and drawing out the life-cycle patterns of SSW

along the age dimension, we show how SSW starts out negative at young ages, increases steadily through retirement, and then gradually decreases as the remaining expected years of life (and thus years of expected benefit receipt) decline at older ages. Age ranges in which multiple cohorts overlap also show us how average SSW is evolving across cohorts, due to a combination of demographic (through spouse/survivor benefits) and lifetime earnings effects.

We begin the life-cycle analysis with means of expected SSW for the 1930 to 1939 through 1970 to 1979 birth cohorts (Figure 7). Each point represents the mean SSW for the indicated birth cohort, plotted at the midpoint of their 10-year age band. Thus, for example, the first marker (blue circle) for the 1940 to 1949 cohort is observed when they are (on average) 50 years old in the 1995 SCF. The subsequent blue markers then follow the 1940 to 1949 cohort forward across survey waves. In 1998, the mid-point of their age range was 53, and increasing by three years until they reached an average age of 74 in 2019. Similarly, the first time we observe the 1970-79 birth cohort is when their midpoint age was 20, in the 1995 SCF. By 2019, the midpoint of their age range was 44. In addition to following the same cohort across survey waves, the cohort-age chart makes it possible to compare two birth cohorts at similar ages but in different survey years. For example, we observe the 1970 to 1979 cohort between ages 20 and 44 in the nine survey waves, and we observe the 1960 to 1969 cohort between ages 30 and 54. In the overlap range between cohorts — ages 30 to 44 in this case — we can investigate how the mean SSW of a given cohort compares to a cohort ahead or behind them.

Figure 7. Mean real expected SSW by birth cohort and age




Notes: Values are per person for respondents and spouse/partners. Projected benefits include own worker and potential spouse and survivor when spouse/partner is observed. Actual benefits used for those already claiming. Taxes are based on own projected earnings.

With those chart-reading principles in mind, we can interpret the estimates in Figure 7. First and foremost, there is a clear life-cycle pattern of SSW accumulation and decumulation. Average expected SSW is negative at younger ages, increases steadily through benefit claiming age, then declines slowly as the number of years of remaining benefit receipt falls with life expectancy. Second, although each cohort has its own internal consistencies as they age through the survey waves, the differences across

cohorts are also noteworthy. For example, the mean SSW estimates for the 1940 to 1949 cohort lie above the 1930 to 1939 values in the age range overlap (ages 60 to 74) but the differences between younger cohorts are less clear or nonexistent. There are possible explanations involving data problems, including the fact that actual benefits may be more underreported at older ages, or the actual and estimated benefits may diverge for other reasons. However, there are also two economic forces that could be causing the convergence in cohort means: real earnings and auxiliary benefits. If lack of earnings growth means that a given cohort has the same average lifetime earnings profile as their predecessors, they will have the same expected SSW at any given age. If a given cohort has more earnings convergence between spouses but the same average earnings, the incremental effects of spouse and survivor benefits are diminished.

Figure 8. Expected SSW relative to income by birth cohort and age




Notes: Values are relative to SCF permanent income measure. Projected benefits include own worker and potential spouse and survivor when spouse/partner is observed. Actual benefits used for those already claiming. Taxes are based on own projected earnings.

The importance of differences in lifetime earnings across birth cohorts can be seen by normalizing estimated SSW by income (Figure 8).²¹ The immediate visual effect is to bring the cohort SSW points into a much clearer life-cycle alignment. This makes sense, because Social Security taxes and benefits are tied to incomes and are fully indexed, thus (other than changes such as the increased Full Retirement Age) real

²¹ The specific income measure used is the SCF “usual” income proxy for normal income, which removes income fluctuations due to transitory income shocks. See Bhutta et al. (2020) for a discussion of the usual income measure.

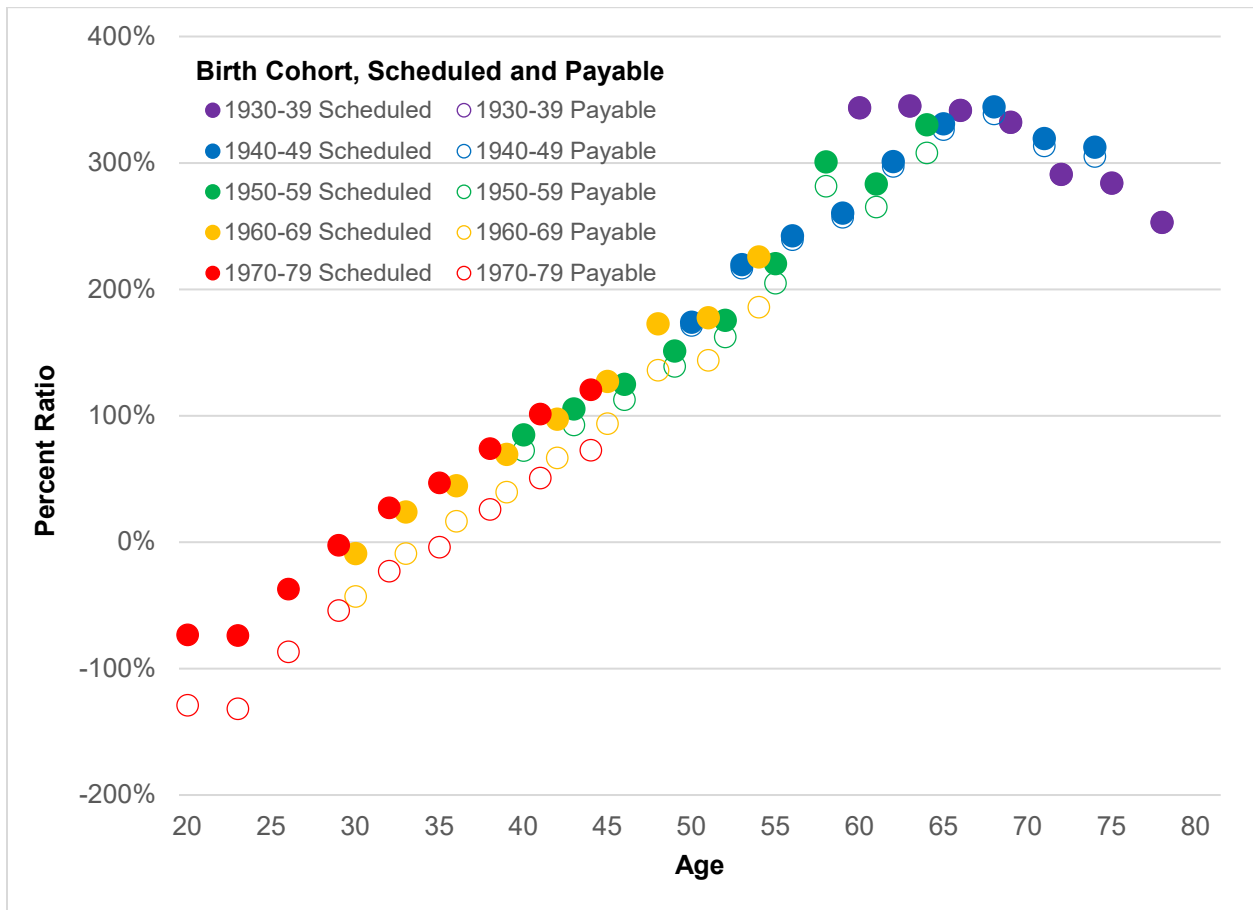
taxes and benefits depend only on real income. The estimated pattern of SSW accumulation relative to income is now very much in line with the theoretical discussion above (Section 2), systematically increasing relative to income through retirement age, before systematically declining at older ages. The remaining differences across cohorts (in the overlap regions) are attributable to factors such as demographics (through spouse and survivor benefits), earnings inequality (through benefit replacement rate differentials), and data (through the assumptions made to construct life-cycle earnings inputs at different ages).

Figure 9. Expected/termination SSW relative to income by birth cohort and age



Notes: Values are relative to SCF permanent income measure. Projected benefits include own worker and potential spouse and survivor when spouse/partner is observed. Actual benefits used for those already claiming. Taxes are based on own projected earnings.

Figure 10. Expected SSW relative to income by birth cohort and age, scheduled and payable



Notes: Values are relative to SCF permanent income measure. Projected benefits include own worker and potential spouse and survivor when spouse/partner is observed. Actual benefits used for those already claiming. Taxes are based on own projected earnings.

Although most of our focus throughout the paper is on expected SSW, we also show the life-cycle means for termination SSW relative to income overlaid on the same chart with the expected SSW estimates (Figure 9). Termination SSW is above expected SSW at younger ages, and the two measures, as expected, converge at older ages. As noted above, under the expected SSW concept, on average, the PDV of benefits reflects the higher benefit the worker will receive because of additional years worked,

but that comes at the cost of additional taxes. Details about the Social Security benefit formula help make it clear why termination SSW is likely so much higher for many participants, especially at younger ages. A young worker with high lifetime earnings will have a low average indexed monthly earnings (AIME) because they have many zeroes in their work history. If they stop working when young, they will be entitled to the benefit replacement rate of a lifetime low earner, which is a higher rate of return. Additional years of work raises their AIME and reduces the rate of return on additional taxes paid.

Our third cohort-age life-cycle SSW chart shows the effects of moving between the scheduled and payable scenarios (Figure 10). Under the payable scenario, we follow recent *Trustees Reports* in assuming that 80% of benefits are payable beginning in 2035. We apply the benefit reduction to all benefits paid from that point forward. The life-cycle chart shows that older cohorts (1930 to 1939 and 1940 to 1949) are somewhere between completely unaffected and only marginally affected, because most of their benefits will have been received before the Social Security trust fund asset reserves are expected to be depleted. Beginning with the 1950 to 1959 cohort — who will be between 76 and 85 in 2035 — the effect of benefit cuts become very noticeable.²²As noted above, the SSW perspective implications for cohort fairness are worse than the 20% benefit cut suggests, because future cohorts will still be paying scheduled taxes. Thus, SSW will fall proportionally more than 20%.

²² It is worth a reminder that the wedge between scheduled and payable here is not being driven by changes in expectations about system finances. Indeed, we use the estimated benefits payable ratio from the most recent *Trustees Reports*, effectively assuming those projections were the same in 1995 as they are today.

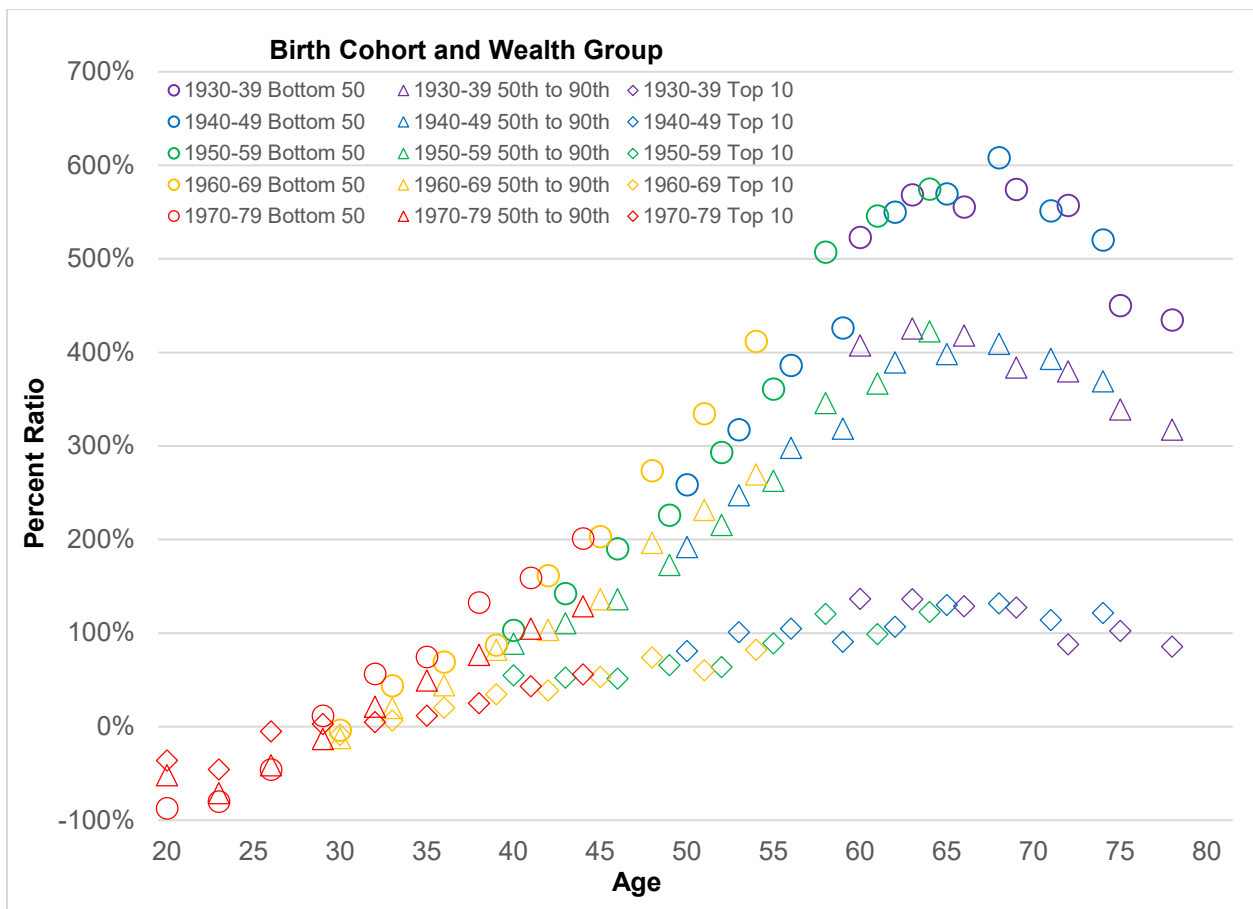
In addition to looking at life-cycle patterns using cohort averages, we can use the pseudo-panel approach to study within cohort patterns (Figure 11). The pseudo-panel approach works for any group decomposition in which the classifier is relatively stable between survey waves. Birth cohort itself is perfectly stable, because the probability of any given individual changing birth year from one survey wave to another is zero. We use the same household wealth classifier developed in Section 4 and used for the age group analysis in the last section to construct within-cohort patterns, which effectively assumes, for example, that a randomly sampled individual in the bottom half of their cohort wealth distribution in one survey year is unlikely to be in a different part of their cohort wealth distribution in a subsequent wave.²³ The approach is still robust if the individuals who move back and forth across wealth groups are similar in terms of the outcome of interest, which in this case is expected SSW.

The general patterns of SSW relative to income are similar across wealth groups, but the important differences show up in the growth of SSW relative to lifetime income during prime age working years. All three wealth groups have little or no SSW at age 30. By age 60 or so, the bottom 50% wealth group has accumulated SSW roughly equal to 600% of their income at age 60. For the 50th to 90th percentile wealth group, accumulated SSW is roughly 400% of income around age 60. Finally, for the top 10% wealth group, the SSW ratio is around 100% of income at age 60. Note that this does not mean that average SSW is greater for low wealth individuals. It is still true that SSW

²³ This principle underscores the desirability of the “person-weighted within age group” approach to classifying observations for distributional analysis discussed in Section 4. Individuals are less likely to systematically change wealth groups and bias the pseudo-panel results if life-cycle and demographic effects are removed.

rises with wealth because wealth rises with income, and SSW rises with income, at least for individuals on the verge of retirement at age 60. Indeed, average SSW peaks at just more than \$150,000 for the bottom 50%, at just over \$200,000 for the 50th to 90th percentile group, and around \$250,000 for the top 10%. The negative relationship in Figure 11 is between SSW *relative to income* and household wealth, and a direct result of the progressive benefit formula.

Figure 11. Mean real expected scheduled SSW relative to income by wealth group



Notes: Values are relative to SCF permanent income measure. Projected benefits include own worker and potential spouse and survivor when spouse/partner is observed. Actual benefits used for those already claiming. Taxes are based on own projected earnings.

The relationship between SSW relative to income and other components of household wealth can also be interpreted in terms of life-cycle saving behavior. We can calculate the implied SSW saving rate by solving for SSW at age 60 using the formula $SSW_{60} = \sum_{a=30}^{60} (s * income_a) * (1 + r)^{60-a}$, where a is age, s is the saving rate, and r is the real interest rate. The SCF pseudo panels suggest that the average real income growth rate for the bottom 50% during our sample period is about 1% per year between ages 30 and 60. The real prime-age income growth rate is about 2.3% for the 50th to 90th percent wealth group, and 5.5 for the top 10% wealth group. Using our baseline discount rate ($r = 2.8\%$), the observed SSW to income ratios translate into annual saving rates of about 15% for the lowest wealth group, 12% for the 50th to 90th percentile wealth group, and 5% for the top 10 wealth group. Conventional savings estimates based on observed household wealth suggest that lower lifetime income families save relatively little (Feiveson and Sabelhaus 2019). In fact, the analysis here shows that low (and even middle) wealth families just happen to be doing most of their retirement saving through the Social Security program.

7. Conclusions

In this paper, we begin with Survey of Consumer Finances (SCF) microdata and create a more comprehensive measure of household wealth that includes the present discounted value of Social Security benefits less taxes, a concept we refer to as Social Security Wealth (SSW). There are four takeaways from this exercise. First, aggregate SSW is quantitatively important when compared to the other components of household wealth. Second, SSW is very skewed toward otherwise low-wealth families at any given age, so adding SSW has a big impact on estimated inequality *levels*. Third, adding SSW to conventional wealth measures does not change perceptions about inequality *trends*, such as the increase in top wealth shares, or the growing wealth gaps within and across age groups. Although gaps between high-wealth and low-wealth families are increasing generally, the gap between old high-wealth and young low-wealth families is exploding, especially if one considers that only 80% of benefits are expected to be payable beginning in 2035. Finally, including SSW helps us better understand typical life-cycle wealth profiles and saving patterns, as SSW rises steadily with age through retirement, then falls gradually with remaining lifespan. More importantly, growth of SSW relative to income — a measure of saving — differs systematically across wealth groups.

The above takeaways are robust to data construction decisions regarding life-cycle earnings and benefits, discounting, and the SSW concept itself, but there are several aspects of our approach that are deserving of further research. On the data front, there is a clear need to test alternative approaches to constructing life-cycle earnings profiles, perhaps multiply-imputing earnings using alternative data sources with longitudinal earnings histories where the donor records match SCF respondents in

terms of characteristics and common labor force and earnings variables. The benefit calculations can be expanded to include other benefit types, including disability and auxiliary benefits based on previous (and perhaps even prospective) marriages. The earnings history imputations and benefit calculator should also be checked against actual respondent-reported benefits in the SCF. In order to fully benchmark the SCF estimates against OCACT values (see Appendix), we will need to add the PDVs for individuals 15 and older who are not represented in the population of SCF respondents and who are either respondents or spouses/partners. Finally, given refined and benchmarked measures of SSW, it will be interesting to bring those measures to bear on explaining macroeconomic outcomes such as consumption and labor supply.

8. 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.
- Catherine, Sylvain, Max Miller, and Natasha Sarin. 2020. "Social Security and Trends in Inequality," Working Paper, Wharton School of The University of Pennsylvania. (February)
- Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The Association Between Income and Life Expectancy in the United States, 2001-2014," *Journal of the American Medical Association*, 315(16):1750-1766. doi:10.1001/jama.2016.4226. (April)
- Coile, Courtney C., and Jonathan Gruber. 2007. "Future Social Security Entitlements and the Retirement Decision," *The Review of Economics and Statistics*, 89(2): 234–246. (May)

- Dudel, Christian, and Mikko Myrskylä. 2017. "Working Life Expectancy at Age 50 in the United States and the Impact of the Great Recession," *Demography*, 54:2101–2123.
- Fang, Chichun, Charles Brown, and David Weir. 2016. "Cohort Changes in Social Security Benefits and Pension Wealth," Working Paper 2016-350, University of Michigan Retirement Research Center. (September)
- Feiveson, Laura, and John Sabelhaus. 2019. "Lifecycle Patterns of Saving and Wealth Accumulation," Finance and Economics Discussion Series 2019-010. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2019.010r1>.
- Goda, Gopi Shah, John B. Shoven, and Sita Nataraj Slavov. 2011. "Implicit Taxes on Work from Social Security and Medicare," in *Tax Policy and the Economy*, 25:69–88, ed. Jeffrey Brown, University of Chicago Press.
- Jacobs, Lindsay, Elizabeth Llanes, Kevin Moore, Jeffrey Thompson, and Alice Henriques Volz. 2019. "Wealth Distribution and Retirement Preparation Among Early Savers," Federal Reserve Board of Governors.
- Leimer, Dean. 2016. "The Legacy Debt Associated with Past Social Security Transfers," *Social Security Bulletin*, 76(3): 1-15.
- 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. 2020. "Social Security Wealth, Inequality, and Lifecycle Savings," National Bureau of Economic Research Working Paper 27110. (May)
- 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.
- Social Security Administration (SSA). 2019. The 2019 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds, available at <https://www.ssa.gov/OACT/TR/2019/>.
- 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

Appendix A. Benchmarking aggregate SSW to OCACT estimates

The SCF-based SSW estimates reported in the text have a published counterpart going back to the 2001 survey (Appendix Table 1). The OCACT benchmarks are a byproduct of disaggregating the “infinite horizon” unfunded liabilities of the Social Security system. The starting point is to compute the discounted present value of all future taxes and all future benefits, take the difference, then add the current trust fund balance to measure the overall aggregate shortfall in present value terms. The decomposition in tables (which are available only since 2001) allocates the PDVs across past, current, and future participants. The groups closest to the SCF population is current participants, which is the population 15 and older in the year for which the calculations are being made.

The SCF population we consider in this paper is a subset of the OCACT current participants group, because we only compute SSW for SCF respondents and spouses/partners, which means we are excluding dependents and other persons 15 and older (the “Non-Primary Economic Unit,” or “NPEU” in SCF parlance) living in SCF households. Adult members of SCF households we miss include children still living with their parents, roommates, parents, or other older relatives living with respondents and spouses/partners. It may be feasible to study those individuals using the rudimentary SCF information that is collected, but for now, it is clear that we expect to calculate PDVs for taxes and benefits that are below OCACT, even if we have the earnings profiles right for respondents and spouse/partners.

In general, the comparison of our estimated PDVs against published OCACT values is reassuring, though the effect of alternative discounting and differences

between the 62 and older and younger than 62 populations are notable and warrant further investigation. On the discounting front, there is little difference between simply using a 2.8% real discount rate and using the year-by-year OCACT discount factors through 2013. Between the 2013 and 2016 surveys, in addition to continued gradual lowering of assumed long-run real discount rates, OCACT moved to an alternative time path for closing the gap between the current (and persistently low) real discount rates and their long-run values, which has the effect of dramatically increasing SSW relative to earlier years. The differences in PDVs between the retirement age and preretirement populations is also notable and warrants further investigation, especially the low PDV of estimated taxes for the 62 and older population.

Appendix Table 1. Survey of Consumer Finances (SCF) and Office of the Chief Actuary (OCACT) present values

	All Current Participants						
	2001	2004	2007	2010	2013	2016	2019
PDV Expected Social Security							
OCACT Published	\$ 10,542	\$ 12,552	\$ 16,265	\$ 19,735	\$ 26,100	\$ 31,400	\$ 37,566
SCF, OCACT Discount Rates	\$ 8,489	\$ 10,327	\$ 13,230	\$ 14,891	\$ 18,561	\$ 27,575	\$ 32,846
SCF, Fixed Real Rate=2.8%	\$ 9,507	\$ 11,529	\$ 14,177	\$ 16,180	\$ 19,382	\$ 21,589	\$ 23,931
= PDV Expected Benefits							
OCACT Published	\$ 23,200	\$ 27,351	\$ 34,257	\$ 40,321	\$ 51,600	\$ 62,000	\$ 72,721
SCF, OCACT Discount Rates	\$ 16,333	\$ 19,012	\$ 23,238	\$ 25,484	\$ 30,273	\$ 42,710	\$ 50,484
SCF, Fixed Real Rate=2.8%	\$ 17,614	\$ 20,549	\$ 24,407	\$ 27,168	\$ 31,139	\$ 34,246	\$ 38,254
- PDV Expected Taxes							
OCACT Published	\$ 12,658	\$ 14,799	\$ 17,992	\$ 20,586	\$ 25,500	\$ 30,600	\$ 35,155
SCF, OCACT Discount Rates	\$ 7,844	\$ 8,685	\$ 10,008	\$ 10,593	\$ 11,712	\$ 15,135	\$ 17,637
SCF, Fixed Real Rate=2.8%	\$ 8,107	\$ 9,021	\$ 10,230	\$ 10,988	\$ 11,757	\$ 12,657	\$ 14,323
	Current Participants, Ages <62						
	2001	2004	2007	2010	2013	2016	2019
PDV Expected Social Security							
OCACT Published	\$ 6,595	\$ 8,030	\$ 10,413	\$ 12,311	\$ 16,000	\$ 19,100	\$ 22,223
SCF, OCACT Discount Rates	\$ 5,117	\$ 6,610	\$ 8,361	\$ 8,815	\$ 10,769	\$ 16,170	\$ 19,180
SCF, Fixed Real Rate=2.8%	\$ 6,041	\$ 7,688	\$ 9,214	\$ 9,900	\$ 11,587	\$ 11,964	\$ 12,722
= PDV Expected Benefits							
OCACT Published	\$ 18,944	\$ 22,418	\$ 27,928	\$ 32,225	\$ 40,600	\$ 48,400	\$ 55,826
SCF, OCACT Discount Rates	\$ 12,926	\$ 15,257	\$ 18,290	\$ 19,300	\$ 22,360	\$ 31,152	\$ 36,596
SCF, Fixed Real Rate=2.8%	\$ 14,113	\$ 16,670	\$ 19,365	\$ 20,779	\$ 23,224	\$ 24,478	\$ 26,839
- PDV Expected Taxes							
OCACT Published	\$ 12,349	\$ 14,388	\$ 17,515	\$ 19,914	\$ 24,600	\$ 29,300	\$ 33,603
SCF, OCACT Discount Rates	\$ 7,809	\$ 8,647	\$ 9,930	\$ 10,485	\$ 11,590	\$ 14,982	\$ 17,415
SCF, Fixed Real Rate=2.8%	\$ 8,072	\$ 8,982	\$ 10,151	\$ 10,878	\$ 11,637	\$ 12,514	\$ 14,117
	Current Participants, Ages 62+						
	2001	2004	2007	2010	2013	2016	2019
PDV Expected Social Security							
OCACT Published	\$ 3,947	\$ 4,522	\$ 5,852	\$ 7,424	\$ 10,100	\$ 12,300	\$ 15,343
SCF, OCACT Discount Rates	\$ 3,372	\$ 3,717	\$ 4,869	\$ 6,076	\$ 7,792	\$ 11,405	\$ 13,666
SCF, Fixed Real Rate=2.8%	\$ 3,466	\$ 3,840	\$ 4,963	\$ 6,280	\$ 7,795	\$ 9,625	\$ 11,209
= PDV Expected Benefits							
OCACT Published	\$ 4,256	\$ 4,933	\$ 6,329	\$ 8,096	\$ 11,000	\$ 13,600	\$ 16,895
SCF, OCACT Discount Rates	\$ 3,407	\$ 3,755	\$ 4,948	\$ 6,185	\$ 7,914	\$ 11,558	\$ 13,888
SCF, Fixed Real Rate=2.8%	\$ 3,501	\$ 3,879	\$ 5,043	\$ 6,389	\$ 7,915	\$ 9,768	\$ 11,415
- PDV Expected Taxes							
OCACT Published	\$ 309	\$ 411	\$ 477	\$ 672	\$ 900	\$ 1,300	\$ 1,552
SCF, OCACT Discount Rates	\$ 35	\$ 38	\$ 79	\$ 108	\$ 122	\$ 153	\$ 222
SCF, Fixed Real Rate=2.8%	\$ 35	\$ 39	\$ 79	\$ 110	\$ 120	\$ 143	\$ 206

Notes: OCACT values for 2013 and 2016 from *Trustees Report* appendix tables, prior years and 2019 are from various tables in the *Financial Report of the US Government*. SCF values based on "expected" SSW concept as described in text.

Appendix B. Details on differential mortality correction

The computation of Social Security and Defined Benefit pension wealth requires conditional survival probabilities for every SCF respondent and spouse/partner from their current age through age 100. The starting point for individual mortality is the average mortality rate by gender, cohort, and age from the Social Security Administration Office of the Chief Actuary website: www.ssa.gov/OACT/HistEst/Death/2019/DeathProbabilities2019.html. The second step is a differential mortality adjustment based on Chetty et al. (2016) that assigns relative mortality by income within each gender and cohort group.

The Chetty et al (2016) analysis of differential mortality uses linked income tax and Social Security death records for the period 2001 through 2014. The specific data file we use has number of deaths and populations by gender, income percentile (1 to 100), age (40 to 76), and year (2001 to 2014). The data are available at <https://healthinequality.org/data/>. There are various ways to process the differential mortality estimates to make them suitable for linking to a data file such as the SCF. We use a method that automatically preserves average mortality within gender, age, and birth cohort groups.

The first step in processing the Chetty et al. (2016) mortality data is to compute death rates for each year, gender, age, and income percentile group. We then average over the 14 years of data, which means we are not attempting to capture and project any relative mortality trends over time. The third step is to compute overall average mortality for each gender and age group, and the fourth step involves dividing percentile group mortality by average mortality to compute *relative* mortality for each income

percentile within the gender and age group. By working with relative mortality rates in this way, we know that the weighted average (across income percentiles) mortality within a given gender and age group will match the average mortality when we merge to a data set with mortality by gender and age across birth cohorts, so long as the matching data has 100 equally weighted income percentile groups. We also further smooth relative mortality in an additional step by regressing relative mortality on a cubic polynomial in income percentile. Those smoothed relative mortality rates are then multiplied by the Social Security cohort by gender and age mortality rates to provide absolute mortality. Values of relative mortality for ages between 76 and 100 are linearly interpolated to converge to one at age 100.

Appendix C. Improving SCF life-cycle earnings imputations using HRS-SSA-linked administrative earnings records

The Survey of Consumer Finances (SCF) labor force and earnings modules have provided the key inputs needed to estimate life-cycle earnings and thus Social Security Wealth (SSW) in this research to date. Our ongoing research is focused on using Health and Retirement Study (HRS) data linked to Social Security Administration (SSA) earnings records to improve the life-cycle earnings imputations in the SCF. The approach involves creating comparable populations in the SCF and linked HRS along two dimensions. The first dimension is basic demographics, and the second is patterns of life-cycle labor force attachment within each of the demographic groups. After the data sets are reconciled, the last step is to explore various ways to use the life-cycle earnings details in the HRS-SSA data to impute year-by-year earnings for SCF respondents.

Linked HRS-SSA administrative earnings file

The SCF data from the current job, labor force history, and work expectations modules used to estimate life-cycle earnings is described above (Main text, Section 3). Before describing how we are using HRS-SSA linked data to improve the SCF life-cycle earnings estimates, we discuss how we are developing the various public-use and restricted data sets available from HRS. Our starting point is the public-use HRS RAND longitudinal file, which pools and links information across survey respondents for HRS interview years 1992 through 2016. The RAND file is the source of the key demographic linking variables.

We merge the HRS RAND file with processed restricted-use SSA administrative earnings data for 1951 through 2012 (using the HRS produced SSA “imputations” file) and then add in subsequent year by year records for those individuals from the detailed earnings file for years 2013 through 2016. The current HRS RAND file has approximately 42,000 records overall, and roughly 24,000 records have links to at least some SSA earnings data. However, many of the 42,000 observations are out-of-scope for our purposes because they were born before our first analysis cohort (1930 to 1939).

Changes in respondent permission requirements for linking SSA earnings records imply that some observations have incomplete earnings — the respondent agreed to have their earnings linked in one survey wave and then changed their mind. With that in mind, we created the concept of “complete earnings” to indicate that the HRS respondent has SSA earnings available through the minimum of (1) their current age, (2) age 61, or (3) their year of death. With that restriction, we have about 14,000 longitudinal observations with complete earnings within our analysis population. The

fraction of observations with complete earnings is just over 50% in the 1930s and 1940s birth cohorts, then drops to around 40% in the 1950s and 1960-64 birth cohorts. There are roughly 5,000 usable observations in the 1930s cohort, but that drops to about 4,000 in the 1940s cohort, about 3,500 in the 1950s cohort, and about 1,700 in the 1960 to 1964 birth cohort group.

Reconciling demographic types across data sets

The choice of demographic variables for disaggregating and aligning the HRS-SSA and SCF populations is based on the tradeoff between variability in earnings levels and life-cycle profiles versus sample sizes across the demographic groups. The four key demographics we are currently working with include (1) birth cohort, (2) gender, (3) education, and (4) occupation. Our on-going work is focused on how to collapse/combine groups to efficiently preserve sample size in one or both data sets being reconciled.

As indicated in the main text, the focus of our output is on 10-year birth cohorts born 1930 to 1939 and 1970 to 1979, but that does not mean we have to reconcile the samples by those same birth years. The HRS-SSA linked sample is generally limited to birth years before 1965 (there are a handful of spouses/partners born in 1965 and later, but those are not useful for our purposes). One key overarching question is the extent to which earnings levels and profiles have shifted across birth cohorts, especially for sub-groups (men versus women, college educated versus others). Our current approach is focused on using three to four birth cohorts to disaggregate the samples at the birth cohort level. Within each cohort, the second (and binary) disaggregation is by gender.

Given those two top-level choices, each of the disaggregated cohort by gender data sets is still well populated, averaging about 1,500 longitudinal observations.

Grouping by education and occupation are also key, because of well-known differences in earnings levels, life-cycle patterns, and variability across those groups after controlling for birth cohort and gender. Our on-going work is focused on balancing lost information (from aggregating) against sampling variability from creating groups that are too small. In terms of education, the key distinction is college educated versus all others. In terms of occupation, the SCF is somewhat limited to begin with as there are six occupational categories in the public-use data: We also distinguish wage earners versus self-employed workers. Thus, we collapse occupation in the SCF and HRS into five comparable groups: (1) managerial and professional, (2) sales and office service, (3) construction, maintenance, production, transportation, and (4) self-employed.

In simplest terms, this means the average 1,500 observations in each birth cohort/gender group are further divided into two education and five occupation groups. The longitudinal sample size is still fairly robust with 150 per birth cohort/gender/education/occupation group on average, but differences in the share of the populations in each demographic group mean that some of the samples are much smaller. Our current focus is on disaggregating in ways that preserve as much of the heterogeneity across groups as possible, by identifying and collapsing along dimensions where there are no statistically discernible differences.

Within-demographics variation in labor force attachment

Disaggregating the HRS-SSA and SCF samples by demographics (birth cohort, gender, education, and occupation) is the first step in reconciling the data sets. However, estimating individual life-cycle earnings involves controlling for the heterogeneity in life-cycle labor supply within those groups. The demographic splits tell us how much the average person of a particular demographic type earned at a particular age, but average earning includes individuals with a range of labor force attachments. This step can best be thought of as assigning work status for every individual at every age, where work status is (1) part time, (2) full time, or (3) not working.

The SCF labor force module metadata (see Section 3 in the main text) already provides what we need to identify the overall degree of labor force attachment for every individual. We know, for example, the number of years the individual has worked full time and part time. Although we do not know at exactly which years (meaning ages) the individual worked full-time, in most cases it is straightforward to make the inference. For example, most 62-year-old men will answer that they have worked full time for something like 40 to 44 years, and the answers are generally consistent with other known life-cycle outcomes such as whether they report having gone to college. Some labor force participation patterns are more complicated to recover, but the various answers respondents provide about their current job or longest career job are generally sufficient to piece together a reasonable life-cycle pattern.

On the HRS-SSA side there is less metadata in one sense because the HRS only asks about the total number of years worked. Part-time and full-time years worked

are not distinguished in the HRS retrospective question. However, there is much more data in another sense because we know the exact earnings at every age from the linked SSA data. We are currently focused on allocating HRS person-level labor force attachment into part- and full-time and reconciling those (within demographic groups) with the SCF metadata. The first step within the HRS-SSA linked file is to reconcile the survey-based number of years worked variable with the count of non-zero earnings years in the SSA records. The second step is to then reconcile total years worked with the sum of part- and full-time years in the SCF by demographic group. The third step is to use the SCF part-time/full-time split and the actual earnings values in the linked HRS-SSA data to reconcile life-cycle labor force attachment measures by age.

Imputing annual earnings over the life cycle

Disaggregating the SCF and HRS-SSA samples by both demographics and patterns of life-cycle labor force attachment is the first step in estimating life-cycle earnings for SCF respondents. The second step is modeling within-group, individual-earnings heterogeneity conditional on the life-cycle labor force attachment measures. One aspect of earnings heterogeneity is how to control for differences between observed SCF earnings (in the year(s) for which we observe SCF earnings) relative to the average within the demographic group at the same point in the life cycle. The earnings heterogeneity adjustment also involves adjusting for the job type at the specific phase of the life cycle, meaning precareer, career, and post-career (or “bridge”) jobs.

The SCF provides up to two distinct observations on individual earnings. If the individual is currently working, the SCF collects earnings details for their current and any second jobs. Regardless of whether the individual is working and has this “main”

job information, they are also asked about their longest lifetime job (or second longest, if their current job is their longest). If that “long” job is longer than five years, the respondent is asked detailed questions about that “long” job, including what they were earning when they left the job.

This combination of questions about “main” and “long” jobs gives us up to two annual earnings observations per individual. In our current SCF life-cycle earnings imputations, those two data points, combined with estimated earnings profiles by gender and education, are used to fill in earnings at every point of the life cycle — both backward and forward relative to the respondent’s current age and age at which they left their “long” job. Those imputations, of course, reflect the labor force attachment variables (not working, part time, and full time).

The same SCF earnings observations are the starting point when introducing heterogeneity using the linked HRS-SSA earnings data, but there are of course more options for exactly how to construct person-level year-by-year earnings imputations. The simplest approach is just to treat the differential between the SCF earnings and the HRS-SSA demographic group average earnings as a fixed effect and preserve that differential at all other points in the life cycle. That approach is similar to what we are currently doing with the SCF data, with the added information that comes from comprehensive life-cycle earnings profiles estimated using HRS-SSA across demographic types. More advanced approaches will involve multiple imputations that capture realistic variability around the life-cycle averages, using a standard permanent versus transitory shock decomposition, where those moments are estimated using the HRS-SSA longitudinal data. Again, as with our current approach, the actual earnings

imputation builds on the individual's labor force attachment by age already inferred from the SCF meta data.

Finally, the imputation involves adjusting for the job type at the specific phase of the life cycle, meaning precareer, career, and post-career (or "bridge") jobs. Unadjusted earnings profiles have the well-known pattern of relatively rapid growth at young ages, followed by a gradual slowing of earnings growth, and eventually a decrease in earnings at older ages. These patterns vary by the demographic variables we are controlling for, but more importantly, the averages reflect moving from precareer (often part-time) to career (usually full-time) jobs, then moving back to post-career jobs (again, often part-time). This distinction matters for our purposes because, for example, an individual who continues working in their career job at older ages generally does not experience an earnings decline and is thus accruing Social Security benefits at a different rate than a similar individual who has moved to a post-career job. Both the SCF and HRS-SSA data sets make it possible to identify these sorts of switches and, thus, we can impute life-cycle earnings more efficiently for individuals based on where they are in terms of life-cycle jobs.