The Evolution of Late-Life Income and Assets: Measurement in IRS Tax Data and Three Household Surveys

James J. Choi
Yale University and NBER

Lucas Goodman
US Department of the Treasury

Justin Katz
Harvard University

David Laibson
Harvard University and NBER

Shanthi Ramnath
Federal Reserve Bank of Chicago

November 23, 2020

Abstract: Using a 5% random sample of administrative IRS tax records covering households born from 1933 to 1952, we evaluate how three widely-used household surveys—the Health and Retirement Study, the Survey of Income and Program Participation, and the Current Population Survey—capture the level of and trends in late-life income and assets. First, relative to the tax data, survey data underestimate total income levels and overestimate declines in income at the median during the initial transition from working life to retirement. Survey estimates of median income at age 73 are lower than tax data estimates by an average of 4.5% in the HRS, 14.2% in the SIPP, and 25.1% in the CPS. Median total income declined from 58 to 68 by an average of only 11.7% in the tax data, compared with 24.4% in HRS, 16.8% in SIPP, and 29.0% in CPS. Second, survey sources overestimate income growth across birth cohorts at older ages but do a better job of capturing these trends at younger ages. Third, lower-income households have not experienced income growth across birth cohorts outside of the Social Security system. Averaging across ages 68 to 74, the 25th percentile income excluding Social Security fell by 16.5% from the 1933 birth cohort to the 1943 birth cohort in the tax data. These declines are larger in the HRS (26.9%) and SIPP (45.5%) and smaller in the CPS (11.1%). The fraction of households in the tax data with no non-Social Security income and no assets at age 72 rose from 18.9% to 20.5% from cohorts born in 1933 to 1945. The fraction of such households is captured well by the HRS and SIPP, but overstated by the CPS.

The research reported herein was performed pursuant to grant RDR18000003 from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA, any agency of the Federal Government, the Federal Reserve System, or NBER. 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.

How well do widely used US household surveys—the Health and Retirement Study (HRS), the Survey of Income and Program Participation (SIPP), and the Current Population Survey (CPS)—capture the level of and trends in late-life income and assets? In this paper, we compare estimates from these surveys to those from a 5% random sample of 2001-2017 IRS tax records covering the 1933-1952 birth cohorts. IRS data contain administrative records on income sources such as pensions and IRAs, which may be severely misreported in household surveys (Bee and Mitchell 2017). IRS data hence offer a unique benchmark to assess bias in survey estimates.

We have three main findings. First, for middle-income households' pre-tax income (excluding capital gains), IRS data show higher levels in early retirement and smaller declines during the initial transition to retirement compared to estimates from survey data. Averaging across the 1933-1939 birth cohorts, survey estimates of median pre-tax income at age 73 are lower than tax data estimates by 4.5% in the HRS, 14.2% in the SIPP, and 25.1% in CPS. Averaging across the 1943-1949 birth cohorts, median pre-tax income declined from ages 58 to 68 by only 11.7% in the tax data, compared with 24.4% in HRS, 16.8% in SIPP, and 29.0% in CPS. Therefore, relying exclusively on survey data to evaluate income levels throughout retirement or trends as households age during early retirement will tend to produce overly pessimistic results about the median.

These patterns are driven by surveys overestimating working-life income and underestimating retirement income. Specifically, wage income at age 58 is overestimated in all three surveys. In contrast, all three surveys tend to significantly underestimate pension income and slightly underestimate Social Security income. HRS estimates of the overall pre-tax income distribution tend to align most closely with the tax data, especially at older ages, but the SIPP often more closely matches the tax data when measuring individual income components. This discrepancy is due to offsetting patterns over- and under-estimation of individual income components in the HRS.

We also compare the distribution of IRA balances in the tax data with levels reported in the HRS and the SIPP. We find that reported levels in the HRS are quite accurate, especially at older ages, while the SIPP significantly underestimates balances.

Second, surveys tend to overestimate income growth across birth cohorts at older ages, but do a better job of capturing these trends at younger ages. Averaging across ages 68-74, median pre-tax income (excluding capital gains) grew 16.3% in the tax data from the 1933 birth cohort to the 1943 birth cohort, compared to an increase of 21.7% in the HRS, 23.5% in the SIPP, and 23.6% in the CPS. On the other hand, averaging across ages 58-67, median pre-tax income (excluding capital gains) grew by -0.2% from the 1944 birth cohort to the 1950 birth cohort, compared to -7.8% in the HRS, -0.7% in the SIPP, and -0.9% in the CPS. Overall, relying exclusively on survey data will produce an overly *optimistic* picture of across-cohort trends in retirement income.

Finally, both IRS and survey data suggest that later-life income outside of the Social Security system has been falling for low-income households with each successive birth cohort. Averaging across ages 68 to 74, the 25th percentile income excluding Social Security fell by 16.5% from the 1933 birth cohort to the 1943 birth cohort in the tax data. Survey data also show declines that are larger in the HRS (26.9%) and SIPP (45.5%) and smaller in the CPS (11.1%). IRS data show that the fraction of tax units at age 72 with no financial assets and no non-Social Security income rose from 18.9% in the 1933 birth cohort to 20.5% in the 1945 birth cohort, the most recent year for which measurements are available. Survey data trends are qualitatively similar. The tax data also reveal substantial geographic heterogeneity in the fraction with no assets and no non-Social Security income—for the 1933 birth cohort at age 72, this fraction ranges from 11.1% at the 10th percentile state (Kansas) to 22.1% at the 90th percentile state (Georgia).

This paper builds on previous work evaluating the accuracy of survey data using administrative records. Compared with papers that explicitly link survey respondents to administrative data (e.g. Bee and Mitchell 2017, Dushi, Iams, and Trenkamp, 2017), we are able to: (i) evaluate trends over a longer time period and in a wider variety of surveys, as we are not restricted by the ability to link to survey data; and (ii) compare population-level trends, as we are not restricted to the subset of linked respondents. Compared with papers that compare aggregate income in household surveys to public administrative aggregates (e.g. Chen, Munnell, and Sanzenbacher, 2018), our access to administrative microdata allows for more precise distributional comparisons of a larger set of income and asset variables.

I. Constructing harmonized datasets in tax and survey sources

A. Data and variable construction

IRS tax data. We construct our IRS sample using tax data from 2001-2017 and the Death Master File from the Social Security Administration (SSA). The SSA data include information on date of birth, sex, and date of death. We draw a 5% random sample of individuals born from 1933-1952 who appear in the tax data at least once during our sample period using the last four digits of a person's Tax Identification Number (TIN),[1] dropping individuals after they die. For state-level analysis, we supplement this sample with a larger draw such that we observe at least 7,000 living people within each birth cohort, state, and year.[2]

We construct income variables using (self-reported) tax return data and (third-party reported) information returns data. All dollar amounts are deflated to 2010 dollars using the Consumer Price Index. To characterize sources of income as households age, we group income into five categories (specifying the form from which the income is reported in parentheses):

(i) Wage income (Form W-2)

(ii) Pension income: the sum of distributions from defined benefit and defined contribution plans (Form 1099-R)[3]

(iii) Social Security income, including OASI and DI (Form SSA-1099)

(iv) Capital income: the sum of taxable and tax-exempt interest income (Form 1040), taxable dividend income (Form 1040), and capital gains (Schedule D)

(v) Other income: the sum of income from rent, royalties, partnerships and S-corporations, estates, and trusts (Schedule E), income from sole-proprietorship income (Schedule C), income from unemployment insurance (Form 1040), and net alimony income (Form 1040)

For the income items derived from Form 1040, we override the Form 1040 amount in certain cases. If an individual is a non-filer or reports zero for a particular component, we typically use the information return amount in its place. Specifically, we use Form 1099-DIV for dividend income, Form 1099-INT for taxable interest income, Form 1099-G for unemployment compensation, and Form 1099-MISC (non-employee compensation) for sole proprietorship income.[4]

[1] Technically, we used a person's masked TIN, which is a deterministic scrambling of their TIN.

[2] State-level results concern the fraction of tax units with no IRA and negligible non-Social Security income. Our target of 7,000 observations per birth cohort and year ensures the standard error on the estimate of this fraction does not exceed 0.6%. In some small states, a 100% pull results in fewer than 7,000 observations.

[3] Pension income includes non-taxable distributions but excludes direct rollovers, Section 1035 exchanges, withdrawals of impermissible contributions, and recharacterized contributions.

[4] The one exception is tax-exempt interest, where we lack relevant information throughout part of our study period. For consistency, we only use self-reported values for tax-exempt interest income. Additionally, we do not attempt to use information returns to calculate Schedule D or Schedule E income.

Based on these income components, we construct several aggregate variables. Pre-tax income sums income in categories (i)-(v). After tax-income is pre-tax income, minus federal tax liability reported on Form 1040 and estimated FICA taxes on wages, plus any earned income tax credit or refundable child tax credits reported on Form 1040. For comparisons with survey data, we construct versions of pre-tax income and after-tax income that exclude capital gains income reported on Schedule D. All total income measures exclude Supplemental Security Income (SSI), private insurance payouts, and SNAP, WIC, TANF, and VA benefits because these are not available in the tax data.

We measure IRA balances held in both traditional and Roth IRAs from an information return (Form 5498).

We classify households as receiving income in a certain category if income in that category exceeds \$500. We classify households as having capital income either if capital income exceeds \$500 or if interest income exceeds \$10,000 times the average effective federal funds rate in the year of measurement. This smooths out variation in interest income on fixed asset holdings due to exogenous changes in the interest rate environment.

All variables are constructed at the tax unit level, summing individuals' income and assets with the income and assets of their spouse. We adjust for household size by dividing income and assets by the square root of the number of people in the tax unit (either 1 or 2, since we do not consider dependents). We are not able to link spouses in years in which a person does not file a return. In such years, we impute spousal links using tax returns from other years when they appear as either a primary or secondary filer on a joint return. Specifically, for a non-filer in year t, we check whether the person appeared on a joint return from $t-4$ through $t+1$, and assign a spouse based on the joint filing closest in time to t, prioritizing status in $t-1$ over status in $t+1$.[5] If the person does not appear on a joint filing during that window, we assume they are single in all years.

Survey data. We prepare versions of the HRS, SIPP, and CPS datasets to align as closely as possible with the tax data. For HRS, we use the RAND HRS 2016 Longitudinal and Detailed Imputation files, RAND HRS Fat Files from 1992-2016, and the HRS public use dataset produced by the University of Michigan. For CPS, we pull the March ASEC supplement from 1992-2018

[5] Data on tax filers runs from 1996-2018, meaning this imputation does not introduce truncation issues at the start or end of the panel.

produced by IPUMS at the University of Minnesota. For SIPP, we merge the 1996, 2001, 2004, 2008, and 2014 panels using files maintained at the NBER. While we focus on people born from 1933-1952 observed from 2001-2017—consistent with the tax data sample—we also compute results for all available cohorts and years to analyze longer-horizon trends.

First, we construct the unit of analysis in the survey data to approximate tax units in the IRS data. We cannot directly identify tax units in the survey data, as we typically do not observe filing status. To maximize comparability, we group two survey respondents into a single economic unit if they are identified in the survey as being married, which typically requires that a married couple is cohabiting at the time of the survey.[6] This grouping will fail to link married couples living separately but filing jointly, who *will* be linked in a single tax unit in the tax data. This grouping will incorrectly link cohabiting married non-filers who did not file jointly in the previous four years or subsequent year, who will *not* be linked in a single tax unit in the tax data. Throughout the paper, we refer to units of analysis in the tax data and survey data as "households." We often identify fewer two-person households in the tax data than in the survey data. For example, for the 1933 birth cohort at age 68, 59.7% of IRS tax units have two persons, versus 65.1% in the HRS, 67.3% in the SIPP, and 70.1% in the CPS. On the other hand, the tax data are more aligned at older ages; for the 1933 birth cohort at age 76, the two-person household share is 54.3% in the IRS data, 58.5% in the HRS, 52.6% in the SIPP, and 60.8% in the CPS.

In the CPS, we lack information on date of birth, so cannot always associate individuals with a unique birth cohort. We do know the exact age of respondents up to age 79.[7] Households are interviewed in March and asked about their income over the previous calendar year; we assign birth cohorts by assuming that respondent age has not changed since December 31 of the income reference year. Therefore, around a quarter of individuals assigned to, for example, the 1935 birth cohort were in fact born in 1936. We drop observations of individuals aged 80 and up.

Next, we construct income and asset variables to match definitions in the tax data. In all sources, we exclude income from SSI, VA pensions, and other government transfers. All survey

[6] There are some exceptions to this in SIPP and HRS due to their longitudinal structure. If members of a married couple are cohabiting when they are initially selected to participate in the survey, but are no longer cohabiting in later waves, SIPP and HRS will follow each member of the original couple. It is possible to link the survey records for such married couples if they stay married even once they are no longer cohabiting. In such circumstances, we do link non-cohabiting married couples in SIPP and HRS.

[7] Ages are reported in bins for 80-85, 85-90, etc.

measures of total income and capital income exclude capital gains and losses,[8] so when making direct comparisons with survey data, we use pre-tax income in the tax data that excludes capital gains and losses. Due to the structure of each survey, some definitions vary slightly as follows:

1. In the CPS and SIPP, pension income includes income from paid-up insurance policies, which are not observed in the tax data or the HRS.[9]
2. For some years, CPS pension income is constructed differently than in other sources. The HRS, SIPP, and tax data separately report pension income from each component—income from defined benefit and defined contribution employer-sponsored pensions; income from IRA withdrawals; and income from annuities. Until 2014, the CPS asked broadly about up to two sources of regular retirement income. Households were not asked about each income source separately, irregular withdrawals from retirement accounts, or income beyond the first two sources mentioned. From 2014 onwards, the CPS asked separate questions about retirement account withdrawals and distributions, consistent with other sources.[10]
3. In the HRS, pension income and other income are constructed differently than in other sources. Regarding pension income, the HRS only asks about total IRA withdrawals since the previous interview—a period of approximately two years. We adjust this amount to reflect withdrawals over a 12-month period, assuming a constant withdrawal rate. Regarding other income, the HRS stopped separately asking about alimony income starting in 2004 (HRS wave 7), and instead allowed respondents to report

[8] The CPS does not ask about income from capital gains and losses. The HRS asks questions that would allow us to calculate changes in the value of real estate and net asset holdings across waves. However, we lack data on initial purchase prices for financial assets, making any measure of capital gains not comparable with the tax data. In some panels, the SIPP asked about capital gains reported on tax forms. We do not use SIPP capital gains variables for two reasons: (i) respondents who report capital losses are not asked the amount of capital losses reported, inconsistent with the tax data; and (ii) the SIPP stopped asking about capital gains income in the 2014 panel, meaning capital gains income is missing from 2013-2017.

[9] Both surveys ask about total income from annuities and/or paid-up life insurance policies in a single question, meaning we cannot separate annuity income from paid-up life insurance policy income.

[10] In 2014, the CPS divided randomly divided the March ASEC Supplement sample into two groups. Three-eighths of the sample received an experimental questionnaire with redesigned pension income questions that asked separately about withdrawals from retirement accounts, and the remaining 5/8 received the old versions of questions. The CPS released separate survey weights to make either sample—when used individually—representative of the full US civilian noninstitutionalized population. We use the appropriately weighted three-eighths sample for 2014. Starting in 2015, the redesigned questionnaire was sent to all households.

alimony income in response to a question asking generally about "other income."[11] Alimony income is potentially significant for single divorced households. To include alimony income while preserving consistency in HRS across years, we include all income reported in response to the generic "other income" questions in our HRS measure of other income in all years.

4. Only the HRS and the SIPP include information on assets. In addition to IRA balances, we construct a more comprehensive variable describing household financial assets by summing balances held in IRAs, defined contribution employer-sponsored pensions, stocks, bonds, and generically specified "other savings."

B. Adjusting estimates to represent the US civilian noninstitutionalized population

The tax data do not represent the same population as each survey. To produce comparable estimates, we target the US civilian noninstitutionalized population in each dataset. This population includes all residents of the 50 states and Washington DC but excludes members of the military[12] and people in institutions such as prisons and nursing homes.

The CPS and the SIPP are already designed to represent the US civilian noninstitutionalized population, and hence do not require any additional restrictions. The HRS initially samples from the US civilian noninstitutionalized population, but continues to interview individuals as they move into nursing homes or move abroad. We therefore exclude observations for years in which HRS participants live outside the US or in nursing homes.

The tax data contain records for any individual who files US taxes or receives a US information return. This *excludes* members of the US civilian noninstitutionalized population who do not have contact with the tax system, but *includes* individuals who are outside of the US civilian noninstitutionalized population and hence not represented in all survey datasets—most importantly, people living abroad and residents of institutions such as nursing homes. The tax data also contain members of the military, who are not included in the survey data if they live in military

[11] Specifically, the HRS asks: "(Other than income you have already told me about,) did you [or your husband/wife/partner] receive any other income in [Last Calendar Year], for example, from private disability insurance payments, consulting fees, rent from your home or second home, odd jobs, and so forth? Do not include financial support from relatives or friends."

[12] Both the SIPP and the CPS include military personnel living in civilian housing units (not in military barracks).

barracks. Such individuals are exceedingly rare among the older population that we study, so we ignore the small discrepancy across datasets potentially induced by such individuals.[13]

We adjust the tax data by dropping observations for years in which members of our tax sample live outside the 50 US states and Washington DC, as identified by having a foreign address on either Form 1040 or Form 1042-S (an information return). We then add individuals who have no contact with the tax system to make the tax data representative of the US resident population. Finally, we adjust our estimates to exclude the institutionalized population.

Adding individuals with no contact with the tax system. Any person who does not appear in the tax data must have zero income and assets. We add enough synthetic zero-income individuals to the dataset such that the number of individuals at each age measured in each year matches estimates of the US resident population published by the Census Bureau.[14]

Adjusting estimates to exclude institutionalized individuals. The tax data include a non-negligible fraction of individuals living in institutions—most commonly, nursing homes—at older ages. However, we are not able to directly identify and remove institutionalized individuals from the tax data. We take a different approach to correct for potential bias introduced by the institutionalized population when estimating distributions of income and assets. First, we construct bounds on the estimates, imposing the extreme assumptions that the institutionalized population lies at either the top or bottom of income or asset distributions. Second, we compute baseline estimates by using estimates from survey data of the institutionalized population's relative position within the income and asset distributions.

To be concrete, we aim to estimate the distribution of income or asset variable Y for the non-institutionalized population (denoted by $I = 0$) in birth cohort τ observed at age a, given by the conditional CDF, $F(y|a, \tau, I = 0) = \Pr(Y \leq y|a, \tau, I = 0)$. We do not directly observe $F(y|a, \tau, I = 0)$. Instead, we observe the conditional CDF for the resident population, $F(y|a, \tau)$, which can be decomposed as follows:

---

[13] In 2018, there were only 3,349 active enlisted members across the entire Department of Defense aged 50-55, the oldest age for which counts are reported (US Department of Defense 2018, Table B-15). This is about 0.03% of the civilian labor force within that age range.

[14] Intercensal estimates for single year of age are anchored on the decennial Census, with annual administrative adjustments for births, deaths, and net migration abroad. There are some age-year cells where the count in the tax data exceeds the Census resident population estimate. We do not adjust the tax data in these instances.

$$F(y|a,\tau) = \Pr(Y \le y|a,\tau, I = 0) \times \Pr(I = 0|a,\tau) + \Pr(Y \le y|a,\tau, I = 1) \times \Pr(I = 1|a,\tau).$$

We estimate $\Pr(I = 0|a,\tau)$ and $\Pr(I = 1|a,\tau)$ using data published by the Census Bureau,[15] so with values for $\Pr(Y \le y|a,\tau, I = 1)$, we can recover the conditional CDF for the non-institutionalized population. To obtain lower bounds for percentiles and levels of receipt, we assume $\Pr(Y \le y|a,\tau, I = 1) = 1$, and to obtain upper bounds, we assume $\Pr(Y \le y|a,\tau, I = 1) = 0$.

For our baseline estimates, we use values for $\Pr(Y \le y|a,\tau, I = 1)$ estimated using data on the nursing home population in the HRS. Measurement error may bias estimates of income and asset levels in the HRS. We therefore proceed under the assumption that, conditional on birth cohort and age, the HRS correctly measures the *relative* position of the institutionalized population within the full resident distribution.[16] For a fine mesh of $n \in (0, 100)$, we compute the probability that an institutionalized HRS individual of age a in birth cohort τ has income lower than the nth percentile of HRS-measured resident income for age a and cohort τ. We do this by regressing a dummy for an institutionalized HRS individual having income less than the nth HRS resident percentile on a cubic polynomial in age and a cubic polynomial in birth cohort.[17] We then assume that an institutionalized individual has this same probability of having less than the nth percentile of income in the tax data.[18]

Appendix Figure A1 shows how in the tax data, for the 1941 birth cohort, the bounds and baseline estimates of the pre-tax income distribution of the resident noninstitutionalized population compare with the pre-tax income distribution of the resident population. Each panel plots the difference in log points between the estimated lower bound, baseline estimate, or upper bound of the resident noninstitutionalized population and the estimates from the resident population. For the 25th, 50th, and 75th percentiles, bounds are quite tight until around age 75—typically much less than ±3%. At later ages, as the institutionalized population grows, bounds become considerably

---

[15] The Census Bureau estimates the annual institutionalized population for each single year of age by anchoring on decennial Census values and adjusting annual estimates using administrative records for net migration into institutions.
[16] Our baseline estimates are robust to any rank-preserving measurement error, as well as age and birth cohort trends in measurement error that may not preserve rank.
[17] We lack sufficient observations in the HRS nursing home population to use a more flexible specification.
[18] The resulting estimates should be weakly monotone: $u < u'$ implies that $\Pr(Y \le F^{-1}(u)|a,\tau, I = 1) \le \Pr(Y \le F^{-1}(u')|a,\tau, I = 1)$. Noise in coefficient estimates means this property may not be satisfied. We enforce monotonicity by re-ranking estimates using the method described in Chernozhukov, Fernandez-Val, and Galichon (2010).

wider. The relative position of baseline estimates indicates that the institutionalized population typically has pre-tax income between the 10th and 50th percentiles of the resident pre-tax income distribution. At the 10th percentile, baseline estimates are close to the lower bound, indicating that the lower bound assumption—that the institutionalized population has income above the 10th resident income percentile—is an accurate approximation. At and above the 25th percentile, baseline estimates tend to hug the upper bound, indicating that the upper bound assumption is a more accurate. Unless otherwise noted, we report baseline estimates in results.

II. Comparing levels and trends as households age

A. Levels and trends as households age in tax data

We first examine how the pre-tax income distribution evolves as households age in the tax data. Figure 1A plots baseline estimates for the 10th, 25th, 50th, 75th, and 90th pre-tax income percentiles[19] at each age for a selection of birth cohorts. Figure 1A also displays baseline estimates for the fraction of households receiving pre-tax income above \$500 (in 2010 dollars). Total income steadily declines with age throughout most of the distribution. At the median, total income drops from around \$45,000 at age 55 to around \$25,000 at age 80; at the 25th percentile, total income drops from around \$22,000 at age 55 to around \$14,000 at age 80. The 10th percentile does not decline monotonically, jumping upwards at the Social Security claiming age around 62-64 before plateauing and slowly declining at later ages. The fraction of households receiving at least \$500 in pre-tax income exhibits this pattern as well.

Our pre-tax total income measure does not measure the levels of after-tax liquid income available for consumption. Tax liability may fall as households age, attenuating the decline in total income. Figure 1B investigates this possibility by plotting the distribution of after-tax income. We continue to see a pattern of declining income with age.

To investigate changes in the distribution of income as households age more systematically, we divide our sample into two groups: (i) households in younger birth cohorts (1943-1949) whom we observe at ages 58-68, and (ii) households in older birth cohorts (1933-1939) whom we observe at ages 68-78. We then estimate the proportional change in each percentile of the pre-tax and after-tax income distributions from age 58 to age 68 for the younger cohort, and

[19] To maintain confidentiality, all "percentiles" in the tax data are constructed by averaging the 10 observations nearest to the exact percentile cuts.

from age 68 to age 78 for the older cohort. Table 1A reports pre-tax income results, along with the across-birth-cohort average of the proportional change at each percentile. From age 58 to age 68, the 25th-90th total income percentiles decline by about 12-13%. The 10th percentile *increases* by about 30%. From ages 68-78, the proportional decline in total income is larger for the lower percentiles than for the upper percentiles. The 90th percentile decreases by about 9%, compared with a 14.5% decrease at the median, a 16.6% decrease at the 25th percentile, and a 35.9% decrease at the 10th percentile. Table 1B reports analogous results for the after-tax income distribution. From 58-68, proportional declines are about half as large compared with the pre-tax income distribution for the 25th-90th percentiles. From ages 68-78, proportional declines are slightly smaller, but similar at and below the median.

B. Comparing levels in the tax data and survey data

Next, we compare these results with estimates from the survey data. Since the survey data typically lack information on capital gains and tax liability, we compare total pre-tax income in each survey with pre-tax income excluding capital gains in the tax data. (For brevity, we will henceforth refer to pre-tax income excluding capital gains as simply "pre-tax income.") To get a preliminary sense for estimates in each source, Figure 2 plots the pre-tax income distribution by age, as estimated by the tax data and survey data, for the 1941 birth cohort. Several patterns emerge. First, a higher fraction of households in the survey data have pre-tax income exceeding \$500 than in the tax data. Levels of receipt in the SIPP and the CPS are closer to the tax data than HRS estimates. This partly accounts for the higher estimates of 10th percentile income in the survey data than in the tax data. Second, from the 25th- 90th percentiles, the SIPP and the CPS tend to underestimate levels of total income compared with the tax data, especially at older ages. Third, at the 25th-90th percentiles, the HRS matches the tax data relatively closely at older ages. For younger ages at the 75th and 90th percentile, the HRS estimates are substantially higher than in the tax data.

Table 2 compares the survey data to the tax data more comprehensively. For each birth cohort and age, we compute the percent difference between percentiles in the tax data and the survey data, and average these percentages across birth cohorts. We also compute the relative difference in the probability of receiving at least \$500 of pre-tax income across datasets. Table 2A focuses on measurements at ages 58, 63, and 68 among younger birth cohorts (1943-1949). For context, we also report the across-cohort average values from each survey. Bracketed values are the minimum and maximum level or percent difference for individual birth cohorts within the 1943-1949 range.

All surveys tend to overestimate income at the 10th percentile relative to the tax data; by age 68, reported values are on average 25.0%, 15.6%, and 20.1% higher in the survey data than in the tax data in the HRS, the SIPP, and the CPS, respectively. At higher percentiles, the HRS tends to overestimate income, while the SIPP and CPS tend to underestimate income.

Comparing the accuracy of each survey at ages 63 and 68, the HRS tends to align most closely with the tax data, except at the 10th percentile, where the SIPP performs better. However, at age 58, the HRS often significantly overestimates pre-tax income relative to the tax data—by an average of 16.3% at the 25th percentile, 15.9% at the 75th percentile, and 17.5% at the 90th percentile. At such points in the distribution, the HRS estimates tend to be less accurate than those in the SIPP and the CPS. For all ages, the SIPP outperforms the CPS at the 25th percentile, but the CPS does better at and above the median.

Table 2B produces similar comparisons of survey and tax data estimates of the pre-tax income distribution at ages 68, 73, and 78, averaging across the group of older birth cohorts (1933-1939). As at younger ages, the survey data report much higher estimates for the 10th percentile than the tax data. Above the 10th percentile, HRS estimates are closest to the tax data, and are often quite accurate in absolute terms. At the 25th percentile, the average percent difference between the tax data and the HRS estimates is 1.5%, -2.5%, and 0.4% at age 68, 73, and 78, respectively; at the median, the average percent difference is -3.8, -4.5%, and -10.7% at age 68, 73, and 78, respectively. Both the SIPP and the CPS tend to underestimate levels of total income. However, the SIPP is uniformly better aligned with the tax data than the CPS. For example, at age 73, the average percent difference between the SIPP and the tax data is -8.9% compared with -13.1% in the CPS; at the median, this difference is -14.2% in the SIPP and -25.1% in the CPS.

Overall, surveys overestimate income at the far left tail. The SIPP and CPS underestimate income for the remainder of the distribution. HRS has a tendency to underestimate income in the remainder of the distribution for older birth cohorts, but a tendency to overestimate income in the remainder of the distribution for younger birth cohorts. Across all birth cohorts, HRS has a slight tendency to underestimate median income.

C. Comparing trends as households age in the tax data and survey data

How do these patterns of under- and over-estimation in survey sources affect survey estimates of declines in income as households age? To investigate, Table 3 compares changes in each percentile of the pre-tax income distribution, as well as changes in the probability of receiving at least $500 of pre-tax income, across different ages for fixed birth cohorts. Panel I reports changes from age 58 to age 68, averaging across the group of younger birth cohorts (1943-1949). Panel II reports changes from age 68 to age 78, averaging across the group of older birth cohorts (1933-1939).

In Panel I, the average proportional declines in the 25th-90th percentiles from age 58 to 68 tend to be larger in the survey data than in the tax data. For instance, the tax data median declines by an average of 11.7%, compared with 24.4% in the HRS, 16.8% in the SIPP, and 29.0% in the CPS. The proportional declines in the SIPP align most closely with the tax data, followed by estimates from the HRS, then the CPS. Panel II shows that survey data estimates of proportional declines from age 68 to age 78 tend to be less uniformly biased upwards—sometimes overestimating declines, and sometimes underestimating declines. Survey estimates at these ages more closely approximate the 12.8% average decline in the tax data median: the HRS estimates a 18.1% decline, the SIPP estimates a 10.2% decline, and the CPS estimates a 16.1% decline.

D. Comparing income in individual categories in the tax data and survey data

The results from Panel I of Table 3 suggest that the survey data often overestimate declines in income during the initial transition from working life to retirement. To produce this pattern, surveys must have either higher levels of working life income than the tax data, lower levels of retirement income than the tax data, or a combination of the two.

We explore these possibilities by measuring distributions separately for each income category in the tax data and survey data. Table 4 reports results. Each row shows estimates for the 10th-90th percentiles of the income distribution in a given category at age 58, 68, or 78, along with the probability of this category of income exceeding $500, averaged over either the younger birth cohorts in Panel I (1933-1939) or the older birth cohorts in Panel II (1943-1949). For each survey source, we additionally report the average percent difference between tax data estimates and survey estimates.[20]

[20] In calculating this average, we omit age-cohort measurements with tax data estimates of zero, for which percent differences are not defined.

We first examine sources of income that are more important while households are still working, namely wage income and "other" income. Other income includes income from self-employment and sole-proprietorships, and hence represents a source of labor income. We find that all surveys have a tendency to overestimate wage income, and the HRS overestimates other income, while the SIPP and the CPS tend to underestimate other income. Overall, the SIPP is more closely aligned with the tax data than the HRS or the CPS within these income categories.

Table 4A reports wage income results. Panel I shows that in the tax data, the fraction of households with wage income exceeding \$500 falls substantially, from an average of 77% at age 58 to an average of 39% at age 68. All survey sources accurately capture these extensive margin levels and trends. The HRS and CPS overestimate wage income percentiles across the distribution at both age 58 and age 68. At age 58, HRS 50th-90th percentile estimates are on average between 10.6% and 15.0% higher than tax data estimates, and corresponding CPS estimates are on average between 14.2% and 21.8% higher than tax data estimates. The SIPP is comparatively more aligned with the tax data; while 25th and 50th percentile estimates are too high at age 58, estimates for the 75th and 90th percentiles are quite accurate both at age 58 and age 68. In Panel II, we see that the fraction of households receiving wage income in the tax data continues to fall to an average of 13% by age 78, a trend accurately replicated in the survey data. The HRS and CPS continue to overestimate the 75th and 90th wage income percentiles at age 68 and the 90th income percentile at age 78, although the HRS is comparatively more accurate. At both age 68 and age 78, the SIPP overestimates wage income percentiles as well, with comparable accuracy to the HRS. The fact that extensive margin estimates are largely accurate, but percentile estimates are often too high, suggests that survey respondents accurately report whether they receive any wage income, but tend to overestimate the amount they receive.

Table 4B reports other-income results. The tax data show that receipt of other income is relatively uncommon—29% on average at age 58, which declines to 15% by age 78. The survey data also exhibit declining levels of receipt, although the HRS estimates are too high and estimates from the SIPP and CPS are too low. Turning to percentiles, HRS estimates at the 75th and 90th percentiles are much higher on average than estimates in the tax data. By contrast, estimates from the SIPP and CPS are much lower than estimates in the tax data. These differences are economically significant. For instance, in Panel I, the HRS estimate of the 90th percentile at age

58 is on average \$32,326, compared with an average estimate of \$7,981 in the SIPP, \$10,352 in the CPS, and \$13,260 in the tax data.

HRS estimates of self-employment income may be more accurate than estimates in the tax data and other surveys. First, there is evidence that households substantially underreport self-employment income to the IRS (Feldman and Slemrod 2007). Second, the HRS questionnaire design makes it more likely to capture self-employment income than other surveys. After asking generally about whether a household earned income from work in the past year, the HRS immediately asks specifically about earnings from self-employment. By contrast, the CPS asks generally about income from "own businesses," and the SIPP asks respondents to classify already-reported earnings as coming from work for an employer or from self-employment.

Next, we study sources of income that are relatively more important in retirement, such as pension income, Social Security income, and capital income. All three surveys typically underestimate income from these sources compared with estimates from the tax data.

Table 4C compares the distribution of pension income in the tax data and survey sources. Unsurprisingly, pension income becomes much more important as households age. Panel I shows that the fraction of households with pension income over \$500 in the tax data grows from an average of 33% at 58 to an average of 62% by 68. Pension income at the 75th percentile grows from an average of \$5,837 at 58 to an average of \$21,672 at 68; at the 90th percentile, the average increases from \$26,745 at 58 to \$42,708 at 68. All survey data sources underestimate both pension income percentiles and probabilities of receipt. The HRS and the SIPP underestimate pension income by a similar amount, while the CPS performs worse. For all surveys, estimates at age 68 are closer to the tax data than estimates at age 58.

Panel II exhibits the continued growth in pension income from age 68 to age 78, albeit less dramatic than the increase from age 58 to age 68. By age 78, an average of 70% of households have pension income over \$500, and the 50th, 75th, and 90th percentiles of the distribution are on average \$7,462, \$22,124, and \$42,567, respectively. Patterns of underreporting are similar to those in Panel I. The HRS and the SIPP suffer from similar levels of underreporting, and the CPS performs worse. Within each survey source, estimates at age 68 tend to be closer to the tax data than estimates at age 78. This suggests that survey reports of pension income become less accurate at older ages, in contrast to results from Panel I, where it appeared that estimates improved with age.

Table 4D reports results for Social Security income. In Panel I, the tax data show that receipt of Social Security income goes from uncommon at age 58, with an average of 19% of households receiving income over \$500, to near-universal by age 68, with 91% of households receiving income. Focusing on results at age 68, once Social Security becomes a meaningful source of income for most households, the SIPP produces estimates most closely aligned with the tax data. Results on the extensive margin are on average nearly identical to the tax data, as are estimates of the 25th-90th percentiles. The HRS follows close behind the SIPP in terms of accuracy. Average levels of receipt are about 4 percentage points higher in the HRS than in the tax data; estimates of the 10th and 25th percentiles are higher on average in the HRS than in the tax data; and HRS estimates for upper percentiles are about 5-6% lower on average than tax data estimates. The CPS performs the worst, underestimating the extensive margin by an average of 6.5%, and underestimating the 25th-90th percentiles by between 5.3% and 16.1% on average.

In Panel II, the relative accuracy of the surveys is more ambiguous. As in Panel I, the HRS overestimates the extensive margin relative to the tax data and on average overestimates the 10th and 25th percentiles at both ages. Estimates at the 50th-90th percentiles are lower than the tax data but closer than equivalent estimates in Panel I. In the SIPP, age 68 estimates are between 6.5% and 10.0% lower on average than in the tax data at the 25th-90th percentiles, considerably less accurate than in Panel I. However, at age 78, estimates are uniformly *higher* than in the tax data by between 1.5% and 10.0% on average. CPS estimates are typically more accurate at both age 68 and age 78 than in Panel I. In fact, CPS estimates at age 78 are, on average, quite aligned with the tax data, especially above the 10th percentile.

Lastly, Table 4E compares the capital income distribution across sources. The tax data show that while around 40-50% of households earn capital income, the amounts received are rarely substantial. Even the average 75th capital income percentile is often below \$1,000. In Panel I, the HRS and the SIPP underestimate and the CPS overestimates percentiles and levels of receipt. Average percent differences between the tax and survey estimates are typically large, but since income levels are low, dollar differences are rarely economically significant. In Panel II, all survey sources tend to underestimate capital income. The CPS and HRS tend to perform best, followed by the SIPP. At the 90th percentile, income levels tend to be larger in the tax data—\$11,543 at age 68 and \$8,775 at age 78 on average—such that large average percent differences, especially in the SIPP, are often economically significant.

E. Comparing IRA balances in the tax data and survey data

Consumption need not track income in retirement, as households can spend down accumulated assets. Therefore, accurately measuring asset holdings during late life is important to understanding overall financial wellbeing. The tax data lack comprehensive asset measures, but do have information on IRA balances. We therefore compare the age- and cohort-specific distribution of IRA balances in the tax data with estimates in the survey data to assess the accuracy of survey-derived asset measures.

Figure 3 shows the age-specific distribution of IRA balances in the tax data, in addition to the fraction of households with positive IRA balances, for select birth cohorts. At most ages, fewer than 50% of households have positive IRA balances; percentiles at and below the median are mostly zero and are omitted from the figure. Several patterns emerge. First, the 75th and 90th percentiles of the IRA balance distribution, and to a lesser extent the fraction of households with positive balances, rises from the mid-50s to age 70. Focusing on the 1947 birth cohort from age 54 to age 70, the 75th percentile grows from around \$22,000 to \$125,000, the 90th percentile grows from around \$93,000 to \$392,000, and the fraction with positive balances rises from 44% to 51%. This increase reflects continued IRA contributions at later ages, asset growth within accounts, and the rolling over of employer-sponsored accounts into IRAs. Second, starting when required distributions begin around age 70, percentiles tend to stay flat with age, while the fraction of households with positive balances begins to decline with age. This suggests that while some households completely exhaust their IRA assets, those with positive balances tend not to spend them down. Finally, amounts held in IRAs at fixed ages has grown in more recent birth cohorts. For instance, at age 70, 75th percentile holdings have increased from about \$46,000 in the 1933 birth cohort to around \$125,000 in the 1947 birth cohort. The growth in the fraction of households with positive IRA balances is more modest, increasing at age 70 from 44% in the 1933 birth cohort to 51% in the 1947 birth cohort.

We next examine how well the survey data measures IRA holdings. Since the CPS does not ask about IRA balances, we restrict our analysis to the HRS and the SIPP. Table 5 reports results using a similar structure to Table 2. For each birth cohort and each age, we estimate the percent difference between estimates in the tax and survey data. Panel I reports averages across younger birth cohorts (1943-1949) at ages 58, 63, and 68, where brackets below each estimate indicate minimum and maximum estimates within individual cohorts. Panel II reports analogous estimates constructed by averaging across older birth cohorts (1933-1939) at ages 68, 73, and 78.

The results show that the HRS is quite well-aligned with the tax data, especially at older ages. In Panel I, the HRS underestimates the fraction of households with positive holdings by an average of at most 5.1%. Estimates of the 75th and 90th percentiles are typically too high, sometimes by large amounts at age 58. In Panel II, HRS estimates are more accurate. On the extensive margin, the average difference between tax and survey data estimates ranges from -2.9% to 0.7%; at the 75th percentile, the average difference ranges from -2.6% to 2.0%; and at the 90th percentile, the average difference ranges from -13.7% to 0.3%. The SIPP performs much worse by comparison. In both panels, the fraction of households with positive IRA balances is underestimated by an average of around 20-25%; both the 75th and 90th percentiles are underestimated by an average of between around 40-70%.

III. Comparing trends in income across birth cohorts

A. Trends in total and after-tax income across birth cohorts using tax data

We start by using the tax data to estimate trends in pre-tax income (including capital gains) and after-tax income across birth cohorts. Specifically, we measure changes at fixed ages from the 1933 cohort to the 1943 cohort, measured at ages 68-74, and from the 1944 cohort to the 1950 cohort, measured at ages 58-67.

Table 6A reports levels and proportional changes at various ages and percentiles across the 1944 to 1950 birth cohorts and the 1933 to 1943 birth cohorts for the pre-tax income distribution. In Panel I, we see that pre-tax income at ages 58-67 has stagnated at the median while falling substantially at the 10th and 25th percentiles, even as it rose at the 75th and 90th percentiles. At the older ages shown in Panel II, we see consistent growth across the entire distribution. From the 1933-1943 birth cohorts, the median increases by an average of 14.1%. Increases in percentiles below the median tend to be slightly lower, while increases in percentiles above the median tend to be slightly higher. Table 6B reports equivalent results for the distribution of after-tax income. Average proportional changes in the after-tax income distribution across cohorts are qualitatively similar but tend to be slightly lower in magnitude. The notable exception is the change in the 10th percentile from the 1944 to 1950 birth cohorts, where the after-tax decline of 20.8% is larger than the before-tax decline of 19.6%.

B. Comparing trends across birth cohorts in the tax data and survey data

We now study whether the survey data accurately capture tax data trends in the pre-tax income (excluding capital gains) distribution across birth cohorts. In each survey source, we compute the proportional change in each percentile of the pre-tax income distribution across birth cohorts at fixed ages. Table 7 shows the results. Panel I reports changes from the 1944 cohort to the 1950 birth cohort, averaging across younger ages (58-67). Panel II reports changes from the 1933 cohort to the 1943 cohort, averaging across older ages (68-74). Average changes in the tax data distribution of pre-tax income excluding capital gains are reported for comparison.

In Panel I, we see that the CPS and SIPP do a good job capturing the trends from the 1944 to 1950 birth cohorts. On the other hand, the HRS overestimates growth at the 10th and 25th percentiles, and underestimates growth at the 50th percentile and above. Panel II shows that for the older cohorts, at the 50th and 75th percentiles, all the survey sources tend to overestimate growth in the distribution of total income. The average proportional increase in the median from the 1933 cohort to the 1943 cohort is 16.3% in the tax data, compared to 21.7% in the HRS, 23.5% in the SIPP, and 23.6% in the CPS. At the 75th percentile, the average proportional increase is 19.9% in the tax data, 23.3% in the HRS, 28.9% in the SIPP, and 27.1% in the CPS. Below the median, results are different for each survey. At the 25th percentile, the HRS underestimates the average increase of 14.6% reported in the tax data, while the increase is overestimated in the SIPP and appropriately estimated in the CPS. All three survey sources underestimate the average increase of 17.5% in the tax data at the 10th percentile.

IV. Stagnating non-Social Security income at and below the median

The results on older ages from the previous section suggest improving retirement preparedness across the income distribution in more recent birth cohorts. In this section, we examine the extent to which this is due to income growth *outside* of the Social Security system.

A. Changes in the distribution of non-Social Security income across birth cohorts

In Table 8, we estimate in the tax data levels and proportional changes in non-Social Security income across birth cohorts for fixed ages. Estimates for the 10th percentile are omitted, as non-Social Security income levels there are almost universally zero. We focus here on results in Panel II, which shows changes from the 1933 birth cohort to the 1943 birth cohort.

At and below the median, the proportional changes in non-Social Security income from the 1933 cohort to the 1943 cohort are much lower than the proportional changes in total pre-tax income. On average, the proportional change at the median is 9.4%, compared with the 14.1% for total pre-tax income shown in Table 6A. At the 25th percentile, non-Social Security income *falls* by 16.5%, compared with an increase of 11.0% for total pre-tax income. The proportional changes in non-Social Security income at the 75th and 90th percentiles are similar to the proportional changes in total pre-tax income. This suggests that the lower income percentiles are becoming more reliant on Social Security income in a way that the upper income percentiles are not.

Table 9 compares these results to survey data trends in the distribution of non-Social Security income, adopting a similar structure to Table 7. In Panel II, the survey sources tend to overestimate growth in non-Social Security income at the median and above. The surveys qualitatively capture the decline in the 25th percentile across cohorts, although the HRS and the SIPP overestimate and the CPS underestimates the magnitude of this decline on average.

B. Changes in the fraction of households with no income and no IRA balances across birth cohorts

The previous results suggest that reliance on Social Security has grown in more recent birth cohorts for low-income households. However, it is possible that households with limited non-Social Security income have substantial other savings available for consumption. While we cannot quantify total available savings in the tax data, we can identify households with no financial assets with relative precision. We directly observe IRA balances, and can detect substantial holdings in remaining accounts through income from interest, dividend, capital gains, and distributions from tax-exempt retirement accounts.[21] We classify households in the tax data as entirely reliant on Social Security if non-Social Security income does not exceed \$500 in 2010 dollars and IRA balances are zero. For these households, income is presumably nearly identical to consumption.

Figure 4A plots the fraction of households in the tax data with no non-Social Security income and no IRA balances for each cohort at selected ages. For each age, the fraction is flat to increasing over time. For example, at age 72, the fraction increases from 18.9% in the 1933 birth

[21] By age 72, all households with positive balances in tax-advantaged employer-sponsored savings accounts such as 401(k)s must take an annual distribution.

cohort to 20.5% in the 1945 birth cohort. Reliance on Social Security also substantially increases with age, from around 13% at age 64 to around 25% at age 80.

Figure 4B focuses on age 72 households, and compares trends in Social Security reliance in the tax data and survey data. Since the HRS and the SIPP have information on assets, we are able to construct estimates of the fraction of households with no non-Social Security income and no IRA balances that are comparable to the tax data. While estimates from the HRS and the SIPP are noisier than the tax data, they show similar levels and trends. We also compute an additional measure that instead excludes households in the HRS or SIPP with total financial assets (including IRA balances) exceeding $3,500.[22] This lowers the fraction of households entirely reliant on Social Security by about 1-3 percentage points in the HRS and about 5 percentage points in the SIPP; however, overall trends across cohorts are similar even in the modified series. The CPS lacks data on IRA balances. For comparability, we compute a version of our Social Security reliance measure in the tax data that does not screen out households with positive IRA balances, and compare this series to the fraction of households in the CPS with non-Social Security income below $500. The CPS significantly overestimates the fraction of households that meet even this less restrictive criterion by about 3-5 percentage points.

Lastly, we examine differences across geography in Social Security reliance. Figure 5 displays the fraction of households with no non-Social Security income and no IRA balances in each US state, estimated from the tax data. Social Security reliance is highest in the Deep South states like Mississippi and Louisiana, and lowest in Upper Midwest states like Wisconsin and Minnesota. Table 10 ranks states by the fraction of households with no non-Social Security income and no IRA balances at age 72 in the 1933 birth cohort, and computes the changes in this fraction from the 1933 to 1943 birth cohorts. There is considerable geographic heterogeneity in Social Security reliance, ranging from to 11.1% at the 10th percentile state (Kansas) to 22.1% at the 90th percentile state (Georgia). Different states have also experienced considerably different trends in Social Security reliance from the 1933 to 1943 birth cohorts, from a 10th percentile change of -1.4 percentage points on a base of 19.0% (North Carolina) to a 90th percentile change of +1.7 percentage points on a base of 11.1% (Kansas). Regressing the change in Social Security reliance

[22] We choose the $3,500 threshold because the average monthly Social Security benefit is about $2,500 for married couples, which households who otherwise have no assets might have in their checking account when interviewed, to which we arbitrarily add $1,000.

at age 72 from the 1933 to 1943 birth cohort on its level in the 1933 birth cohort yields a coefficient of -0.15 ($p < 0.001$), indicating geographic convergence over time.

V. Conclusion

The tax data show considerable declines in income across the distribution as households age. However, relying on survey data alone will overstate this decline at the median during the initial transition to retirement, leading to overly dire assessments of middle-income household retirement preparedness. By contrast, while the tax data show income growth across birth cohorts, growth at older ages across cohorts is overstated in survey sources. Leaning on survey data will therefore present an excessively optimistic picture of cohort trends in retirement preparedness. All sources indicate declining income outside of the Social Security system across birth cohorts for low-income households, as well as an increasing share of households who are entirely reliant on Social Security to finance their consumption.

References

Bee, Adam and Joshua Mitchell, 2017. "Do Older Americans Have More Income Than We Think?" SESHD Working Paper #2017-39.

Chen, Anqi, Alicia H. Munnell, and Geoffrey T. Sanzenbacher, 2018. "How Much Income Do Retirees Actually Have? Evaluating the Evidence from Five National Datasets." CRR Working Paper #2018-14.

Chernozhukov, Victor, Iván Fernández-Val, and Alfred Galichon, 2010. "Quantile and Probability Curves without Crossing." *Econometrica* 78(3): 1093-125. Available at: www.jstor.org/stable/40664520.

Dushi, Irena, Howard M. Iams, and Brad Trenkamp, 2017. "The Importance of Social Security Benefits to the Income of the Aged Population." *Social Security Bulletin* 77(2).

Feldman, Naomi E. and Joel Slemrod, 2007. "Estimating tax noncompliance with evidence from unaudited tax returns." *The Economic Journal*, 117(518), 327-352. Available at: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0297.2007.02020.x

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner, 2017. "Lifetime Incomes in the United States over Six Decades." National Bureau of Economic Research Working Paper #23371. Available at: https://www.nber.org/papers/w23371.pdf

US Department of Defense, 2018. "Population Representation in the Military Services: Appendix B tables." Available at: https://www.cna.org/pop-rep/2018/appendixb/appendixb.pdf.