

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Lifecycle Patterns of Saving and Wealth Accumulation

Laura Feiveson and John Sabelhaus

2019-010

Please cite this paper as:

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.2019.010>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Lifecycle Patterns of Saving and Wealth Accumulation

Laura Feiveson
John Sabelhaus

Abstract

Empirical analysis of U.S. income, saving and wealth dynamics is constrained by a lack of high-quality and comprehensive household-level panel data. This paper uses a pseudo-panel approach, tracking *types* of agents by birth cohort and across time through a series of cross-section snapshots synthesized with macro aggregates. The key micro source data is the Survey of Consumer Finances (SCF), which captures the top of the wealth distribution by sampling from administrative records. The SCF has the detailed balance sheet components, incomes, and interfamily transfers needed to use both sides of the intertemporal budget constraint and thus solve for saving and consumption. The results here are consistent with recent papers based on individual panel data from countries with administrative registries, and highlights the different roles of saving, capital gains, and interfamily transfers in wealth change over the lifecycle and across permanent income groups.

Keywords: Household income, consumption, saving, wealth

JEL Codes: D14, H55, J32

Laura Feiveson is Principal Economist and John Sabelhaus Assistant Director, both in the Division of Research and Statistics, Board of Governors of the Federal Reserve System. We are very grateful to Marco Angrisani and Kevin Moore for detailed comments on an earlier draft, and to Elizabeth Holmquist for assistance with the Financial Accounts balance sheet and capital gains data. We also received very helpful feedback and suggestions during presentations at the August 2017 IARIW meetings in Copenhagen, the Federal Reserve Board, and the Washington Center for Equitable Growth. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. Emails: Laura.J.Feiveson@frb.gov, john.sabelhaus@frb.gov.

1. Introduction

Measuring variation in the joint distribution of income, consumption, and wealth over the lifecycle and across different types of consumers is key for addressing two overarching questions in economics. First, rising wealth inequality has led to increased theoretical and empirical work exploring the role of income and saving dynamics in explaining wealth concentration, with various emergent explanations. Some explanations rely on differences in characteristics like patience or individual ability, while other explanations focus on factors such as heterogeneity in labor incomes or returns to capital. Second, there is great interest in the comovement of consumption with income and wealth in response to new information at business cycle frequencies. In particular, consumption responses that are linear in wealth and income—as they are in many workhorse macroeconomic models—will yield changes in output little affected by wealth and income inequality.¹

Studying these wealth concentration and business cycle questions requires a particular type of data that is sorely missing for the U.S. economy. The data that economists would like to have for studying such questions is a large representative panel with well-measured household-level data on incomes, consumption (or saving), and wealth. Such data (or reasonably close approximations) do exist for administrative “registry” countries such as Sweden and Norway, but they are not available for the U.S. economy. Some available U.S. data sets each have key pieces of the overall puzzle, but no one data set has all of the pieces in one place. As such, the answers provided to the two overarching questions above are generally very dependent on which of the incomplete data sets are used, and how.

The main contribution of this paper is to synthesize available U.S. micro and macro data in order to recover the joint distribution of income, consumption, and wealth across groups at lifecycle frequencies. The empirical framework is a pseudo-panel, which means we are tracking *types* of agents over the lifecycle and across time through a series of cross-section snapshots. The key micro source data is from the triennial Survey of Consumer Finances (SCF) for 1995 through 2016. The SCF captures the top of the wealth distribution using a sampling and validation approach based on administrative data.² The SCF also includes direct estimates of disaggregated balance

¹ See, for instance, Krusell and Smith (1998).

² For a description of the latest SCF results and a discussion of the administrative data sampling and validation, see Bricker et al. (2017).

sheet components and the capital incomes associated with each type of wealth, the measures of interfamily transfers needed to complete the intertemporal budget constraint, labor incomes, and key demographic variables. We synthesize the survey snapshots with detailed macro income and wealth time-series, and thus we are able to benchmark the joint distributions of income, saving, and wealth over the two decades (and seven three-year sub-periods) spanned by the 1995 through 2016 SCF data sets.

Consistent with recent studies using administrative registries for other industrial economies, the pseudo-panel wealth change accounting framework presented here focuses attention on the role of asset prices and heterogeneity in rates of return to capital when considering differences in saving over the lifecycle and across time. For example, Bach, Calvet, and Sodini (2018) and Fagerang, Holm, Moll, and Natvik (2018) show that the accounting treatment and estimates of the capital gains component of wealth change is key for interpreting the extent to which differences in savings behavior per se versus heterogeneity in (say) income processes is the key to understanding wealth inequality.³ We are able to show the same basic relationships at the agent-type level in the U.S. using the pseudo-panel approach. In addition, the fact that we observe capital income and wealth for the same households allows us to directly test the assumptions required to solve for saving across capitalized income fractile groups, as in Saez and Zucman (2016).⁴

The first important data innovation required to build the pseudo-panel is synthesizing the micro and macro data for the various intertemporal budget constraint components, which makes it possible to tie the results back to the macroeconomic aggregates and distributional outcomes of interest. We show that the SCF micro data generally line up very well with comparable National Income and Product Account (NIPA) and Financial Account (FA) income and wealth aggregates, so for most income and wealth components we can simply use proportional scaling to reproduce the aggregate intertemporal budget constraints precisely. There are three wealth components—

³ Baker et al. (2018) consider how measurement error in balance sheet components flows through to error in consumption (or saving) using the intertemporal budget constraint approach. Those sorts of errors are relevant for both the registry papers and our pseudo-panel approach. The authors show that the errors are on average small and centered around zero, but they do vary with income and over the business cycle.

⁴ The Saez and Zucman (2016) capitalized income approach to measuring wealth concentration is sensitive to heterogeneity in the rate of return to capital, as explained by Kopczuk (2015), Bricker, et al. (2016), and Bricker, et al. (2018). For the purposes of measuring saving, the key point is that the bias from assuming homogeneous returns in the capitalization model maps directly into biased saving estimates. There are no independent estimates of wealth and income with which to properly separate saving out of income from capital gains, even in the absence of movement across wealth fractiles.

owner occupied housing, non-corporate businesses, and vehicles—for which the aggregates are not easily observed using available administrative or market data, and for which SCF respondents (in aggregate) report higher market values. We interpret the differences between the aggregated micro values and published macro as disagreement between government statisticians and SCF respondents about cumulated capital gains on those assets.⁵ Thus, in our decomposition of wealth change, saving summed across agent types matches published aggregates, while capital gains (on housing, owned businesses, and vehicles) are slightly higher.

A second important data innovation here is explicit accounting for interfamily transfers in the intertemporal budget constraint, including both bequests/inheritances at death and *inter vivos* transfers. The SCF includes respondent-reported values for inheritances received, and for *inter vivos* transfers made and received. We complete the between-agent type interfamily transfer flows by estimating bequests made using a model of differential mortality applied to beginning of period wealth holdings. The simulated bequests are validated by showing that the distribution of estimated bequests made lines up very well with the distribution of reported inheritances received. In the empirical work, we show that accounting for the heterogeneity in transfers made and received is important for the decomposition of wealth change into component sources at various points in the lifecycle.

The lifecycle patterns of wealth accumulation that emerge from the pseudo-panel disaggregation provide new insights about heterogeneity in U.S. saving and wealth accumulation. We focus on decomposing the change in wealth at every age and for various agent types into three components: conventionally measured (NIPA and FA concept) saving, capital gains, and net interfamily transfers received. Similar to individual-level panel data from economies with administrative registries, the pseudo-panel shows the importance of capital gains in accounting for wealth change over the lifecycle, especially for the highest permanent income group. Saving and net interfamily transfers both play important roles in determining wealth change at various points

⁵ It may seem obvious that the published macro aggregates are closer to the truth than the aggregated micro data, and indeed much of the work in this paper and elsewhere that involves synthesizing micro and macro data makes that assumption. However, it is important to remember that the published government aggregates are themselves estimates, and, for example, Gallin et. al (2018) explains Federal Reserve methodology for estimating FA housing values that closes much of the historical gap between FA and SCF. This is empirically important because Bricker et al. (2016) show that some of the divergence between SCF and capitalized income wealth concentration (as reported by Saez and Zucman (2016)) is attributable to differences in aggregate home values based on the old FA methodology.

in the lifecycle, but the patterns also clearly differ across the permanent income measure we use to distinguish agent types.

There are two different ways to think about saving and wealth change in the comprehensive intertemporal budget framework. The ratio of saving to disposable income is our primary measure of the saving rate, because it is the same concept as the personal saving rate in the NIPA and FA, and thus sums over individuals to match the aggregates. In contrast to the sorts of conceptually-inconsistent saving rates that have been measured using cash-flow concepts in available micro data, our pseudo-panel saving shows a clear hump shape over the life cycle, turning negative between ages 50 and 60.⁶ The second way to think about saving is to measure the fraction of resources that flow to the individual not consumed in the current year, where resources include disposable income, interfamily transfers, and capital gains. The second measure helps make it clear why wealth does not decline at older ages: capital gains and net transfers received by surviving agents are more than enough to offset negative saving.

The decomposition of wealth change at various lifecycle stages is also instructive for understanding the joint distribution of income, consumption, and wealth across various agent types. Low permanent income agents have very low savings during their working years, which is unsurprising in hindsight given the low levels of observed wealth for those agent types at any point in the lifecycle. Indeed, the wealth owned by lower-income agents is mostly in the form of housing, and most of the growth in that wealth in the past two decades is because of house price appreciation. The highest permanent income group does exhibit the highest saving (relative to income) at younger ages, roughly double that of the middle income group. However, negative saving at older ages holds for all agent types, and the growing ratio of capital gains on accumulated wealth to income by age is key to understanding why the wealth of the highest permanent income types (relative to income) grows over the entire lifecycle.

This paper contributes directly to the empirical literature on wealth inequality dynamics. The theory laying out the candidate explanations for wealth concentration (above and beyond labor income concentration) is well described by Gabaix et al. (2016), Benhabib et al. (2015, 2017), and Benhabib and Bisson (2018). However, there are open questions about how any given combination

⁶ The conceptual inconsistencies in cash flow saving estimates are mostly due to the treatment of retirement income. Pension payments and withdrawals from IRAs and 401(k) accounts are not part of (NIPA consistent) income, because they represent the drawing down of an existing asset.

of income processes and heterogeneity across agents come together to generate the observed skewness in wealth holdings. Some models, dating back to Krusell and Smith (1998) but as recently as Carroll et al (2017), rely on heterogeneity in discount rates or direct preferences for current versus future consumption in order to generate realistic wealth distributions. Some direct empirical analysis, such as Fagerang, et al. (2016), finds that heterogeneity in the rate of return to capital is a key explanation for deviations from the predictions of Bewley-type models.⁷ Some models such as Casteneda et al. (2003), De Nardi et al. (2016), De Nardi and Fella (2017) focus on non-standard stochastic labor income processes to solve the wealth concentration mystery. Although we find strong evidence of heterogeneity in savings behavior, our results are consistent with the idea that behavior relative to conventionally measured income will never fully explain wealth concentration, because the fraction of wealth change explained by saving out of conventionally measured income is a relatively small component of wealth change. Furthermore, since gains are such an important factor in wealth accumulation, it is imperative that we study more the reasons that drive consumers to choose one type of an asset over another.

The results here are also informative for the more general empirical literature on levels and trends in inequality, as captured by different data sets and for different concepts. The available U.S. micro-level data has provided a wide range of estimates for levels and trends in inequality for income, consumption, and wealth. Some of the differences in levels and trends are to be expected, because theory suggests (for example) that consumption should be more equally distributed than income and wealth due to consumption smoothing and insurance across families. However, some of the differences are due to the sorts of population coverage, conceptual, and measurement problems described by Attanasio and Pistaferri (2016).⁸ The focus in this paper is on using the identities that link the various concepts together at the micro level, and on bringing to bear different types of micro and macro data. By focusing on the complete joint distributions and the relationship between micro and macro variables, we improve understanding about the relationship between income, consumption, and wealth inequality.⁹

⁷ Dynan, Skinner, and Zeldes (2004) also find that savings rises with lifetime income, but reject the idea that those patterns are explained by heterogeneity in rates of time preference.

⁸ Bosworth et al. (1991) is an early example of survey-based attempts to measure saving by differencing reported income and consumption.

⁹ In related work, Fisher et al. (2016a, 2016b) also look at the joint distribution of income, consumption, and wealth using various survey data sets, including the SCF, but they do not focus on the household budget identity that ties the concepts together.

A final contribution of the paper is improving our understanding of key empirical joint distributions that are currently influencing economic policy and forecasting. Disaggregated data on income, consumption, and wealth across agent types has been used to gauge differences in behavior at business cycle frequencies. The pseudo-panel data generated here can in principle be used to inform those same questions, which in turn will help us understand and affect macro outcomes by incorporating the heterogeneity in circumstances and/or behavior over time. For example, a great deal of attention has been paid to the borrowing and spending behavior of different types of agents during the U.S. housing boom, and how spending behavior changed in the subsequent bust. In particular, Mian and Sufi (2011) argue that the availability of credit to lower-income households was a substantial contributor to the boom and bust. The pseudo-panel approach here can be used to investigate differences in borrowing and spending before, during, and after the financial crisis. Indeed, previous work by Devlin-Foltz and Sabelhaus (2016) using the same SCF data used here provides evidence against simple stories about credit availability and mortgage default across agent types.

The rest of the paper is organized as follows. In Section 2 we introduce our intertemporal budget constraint empirical and accounting framework, focusing on the micro/macro data synthesis and interfamily transfers needed to track wealth changes across cohort and agent-type groups. In Section 3 we describe our pseudo-panel methodology for disaggregating wealth change across cohorts and agent types using the synthesized micro/macro data, which involves, among other things, careful tracking of births and deaths in the context of the cross-section surveys. In Section 4 we show the point estimates of per-capita wealth change components, income, and consumption for each birth cohort and across the three-year sub-periods in our samples. Arraying the point estimates along the age dimension provides the first view of the lifecycle patterns we are trying to estimate, and fitting a smoothed line through the point estimates shows the patterns even more clearly. In Section 5, we show smoothed lifecycle wealth change decomposition across permanent income groups, linking groups across the cross-sections by using relative rankings of permanent income within cohorts. Section 6 concludes.

2. The Intertemporal Budget Constraint in Micro and Macro Data

The textbook household intertemporal budget constraint is the starting point for measuring saving and wealth dynamics. The budget constraint links wealth change on the left hand side to saving—disposable income minus consumption—on the right hand side. The goal of this paper is to disaggregate the sources of household wealth change across well-defined agent types, so establishing the conceptual and empirical relationship between the micro and macro data is a crucial first step.¹⁰

Saving in the NIPA and FA

The most widely referenced measure of aggregate household saving is based on the right-hand side of the intertemporal budget constraint, as in the National Income and Product Accounts (NIPA).¹¹ In very broad terms, the concept of saving (S_t) in the NIPA is just disposable income (Y_t) minus consumption (C_t):

$$S_t = Y_t - C_t$$

The most important thing to note from a budget identity perspective is that the NIPA concept of saving does not include capital gains, which we will show is a key driver of wealth change over the lifecycle and across time. The decision not to include capital gains derives from the idea that NIPA seeks to quantify the incomes derived from current production, not the change in wealth.

The Financial Accounts (FA) derivation of aggregate household saving begins with the left-hand side of the budget constraint, which is the change in wealth ($W_t - W_{t-1}$).¹² The household sector of the FA focuses on quantifying the balance sheet position (net worth) of households at any given point in time, and it is straight-forward to difference the point estimates to solve for change in net worth over time. However, in order to conceptually match NIPA saving, only the component of net worth change attributable to saving out of current production is counted.¹³ In FA parlance, it is the “net investment” in assets and “net change” in liabilities that is conceptually consistent with NIPA saving (S_t). The residual component of wealth change

¹⁰ See Online Appendix 1 for a detailed discussion of the adjustments made to the NIPA, FA, and SCF data to create the aligned data sets described in this section.

¹¹ See www.bea.gov/iTable/index_nipa.cfm.

¹² The FA data is described in the Federal Reserve’s Z1 release, see www.federalreserve.gov/releases/z1/current/.

¹³ Gale and Sabelhaus (1999) provide more details and a historical perspective on the theoretical and empirical relationship between FA and NIPA aggregate saving rates.

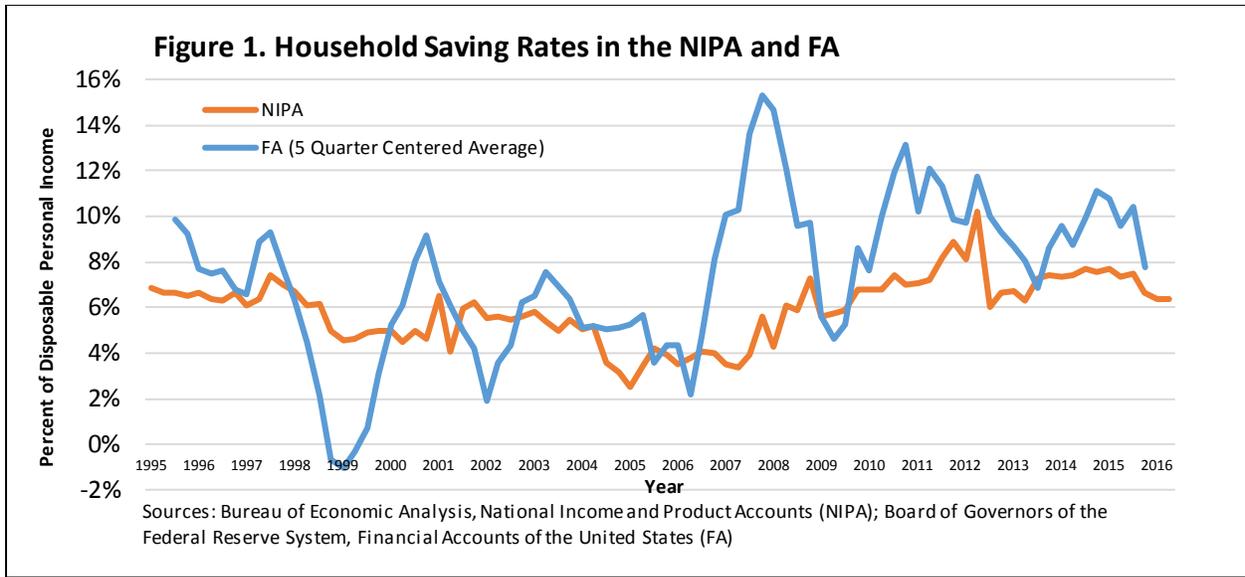
is capital gains, which, in the language of the FA, is “holding gains” on existing assets (G_t). The basic FA wealth change identity is thus:

$$W_t = W_{t-1} + S_t + G_t$$

We can rewrite the identities for change in wealth and flow saving in the form of the usual intertemporal household budget constraint:

$$W_t - W_{t-1} - G_t = Y_t - C_t$$

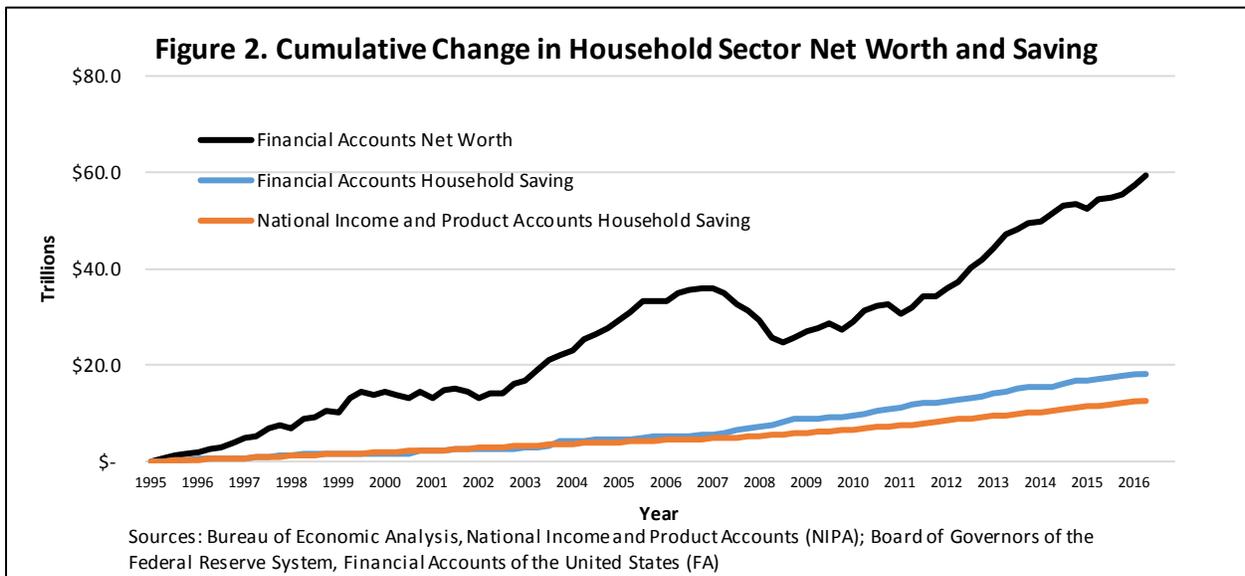
Note, however, that creating a concept of saving that counts holding gains as a component of income (realized holding gains are part of income under the income tax, for example) simply involves moving all or some of G_t to the other side of the identity.



Although the household budget constraint is an identity in principle, even conceptually reconciled NIPA and FA household saving estimates diverge in practice.¹⁴ In general, the conceptually-equivalent FA saving rate fluctuates more than its NIPA counterpart (figure 1). Both measures show that savings, has been on average, about 6 percent of disposable income over the past two decades. Also, both series show the same trend decline in saving rates between the mid-1990s and mid-2000s, but the FA decline is more dramatic, both starting at a higher

¹⁴ Financial Accounts Table F.6 provides the reconciliation between NIPA and FA saving needed to produce this figure. The largest alignment adjustment is removing investment in consumer durables from the FA measures.

level and ending slightly lower. The increase in FA saving post financial crisis is also somewhat more dramatic, rising above the relatively higher levels observed in the mid-1990s.

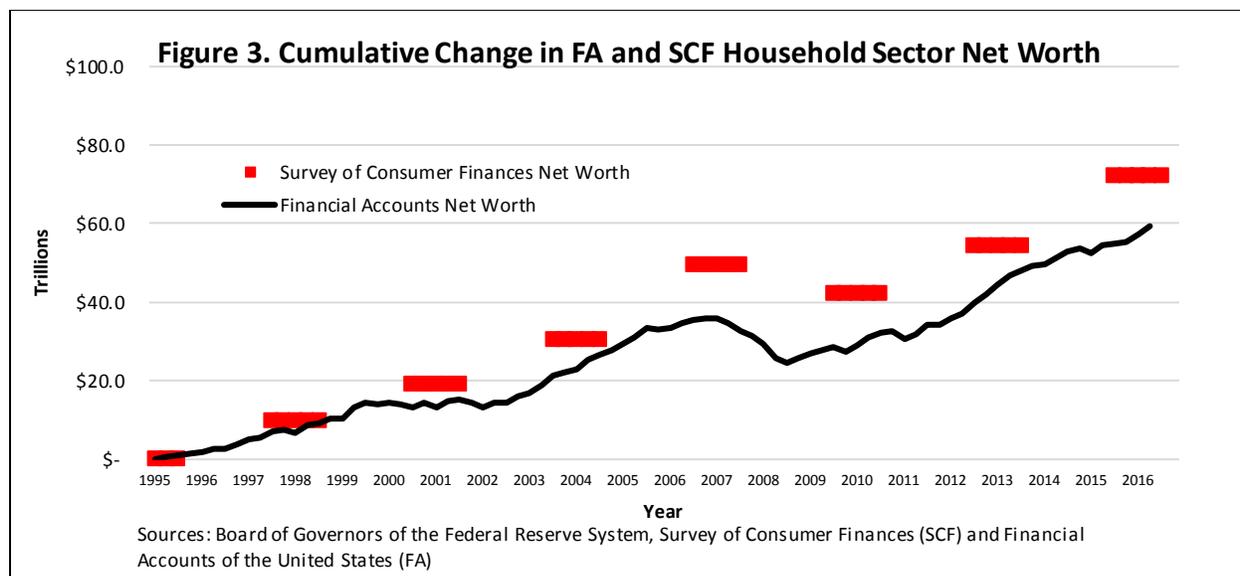


Disentangling Saving from Capital Gains

The concept of saving in the NIPA (which is conceptually the same as net investment less net borrowing in the FA) does not include capital gains. Some perspective on the saving component of wealth change is provided by considering how cumulated flow saving compares to aggregate wealth change over time (figure 2). The chart shows three measures of cumulated wealth change and saving over the period 1995Q1 through 2016Q4. The top (blue) line is the cumulative change in FA household sector net worth, which is almost \$60 trillion for the past two decades. The bottom (red dotted line) is cumulated NIPA personal saving, which is about \$13 trillion over the same period. Thus, saving accounts for less than 25 percent of household wealth change during this period, which suggests capital gains accounts for more than 75 percent of the total. Using the alternative FA saving measure changes the decomposition only slightly, because cumulated saving using that measure is close to \$20 trillion, or just under one-third of total wealth change.¹⁵

¹⁵ The decomposition of wealth change in figure 2 captures corporate retained earnings through capital gains, not saving per se. Obviously retained earnings are a form of saving in a comprehensive private saving measure, but from the perspective of households retained earnings shows up as changes in equity prices.

The FA and NIPA data show that most of aggregate household sector wealth change is accounted for by capital gains, and not by conventionally measured saving. That same relationship has to hold in the aggregated micro data as well, but it does not mean that gains dominate wealth change across all agent types and at all points in the lifecycle. Indeed, to the extent that particular types of agents at particular points in the lifecycle are acquiring net assets, other types of agents at other points in the lifecycle may have an even higher ratio of capital gains to saving. In order to use the micro data to disaggregate wealth change across agent types and lifecycle stages, we first must align aggregated household sector balance sheets in the micro and macro data.



Synthesizing Micro and Macro Balance Sheets

The methodology for collecting micro and macro data on household sector wealth are very different, and even on a conceptually adjusted basis, there are residual differences in aggregated totals.¹⁶ Household sector net worth in the SCF micro data grew much faster than the FA published aggregate over the 1995 through 2016 period (figure 3). While FA aggregate household sector net worth (the blue line, also from figure 2) grew nearly \$60 trillion over the past two decades, the SCF (marked by the red squares spanning each SCF field period) grew

¹⁶ See online appendix 1 for a detailed discussion of the steps taken to align SCF and FA balance sheet components. That appendix is largely based on the work of Dettling et al. (2015), but see also Bricker et al. (2016).

nearly \$80 trillion.¹⁷ Given that the SCF is a survey with sampling and measurement variability, other research has suggested that the SCF is not capturing the value of key balance sheet components properly, and the solution is to benchmark the SCF values (using proportional scaling) to the published FA aggregates.¹⁸

Table 1. Household Net Worth in the Financial Accounts and Survey of Consumer Finances			
Balance Sheet Category	1995 Survey of Consumer Finances	Financial Accounts	
		1995 Q1	1996 Q1
Financial Assets	\$ 16.2	\$ 16.7	\$ 18.7
+ Real Estate	\$ 8.3	\$ 7.8	\$ 8.2
+ Noncorp Business	\$ 4.9	\$ 3.4	\$ 3.6
+ Vehicles	\$ 1.2	\$ 0.9	\$ 0.9
- Liabilities	\$ (3.7)	\$ (4.2)	\$ (4.6)
= Net Worth	\$ 27.0	\$ 24.6	\$ 26.8

Balance Sheet Category	2016 Survey of Consumer Finances	Financial Accounts	
		2016 Q1	2017 Q1
Financial Assets	\$ 63.4	\$ 59.1	\$ 63.9
+ Real Estate	\$ 28.8	\$ 21.9	\$ 23.4
+ Noncorp Business	\$ 21.9	\$ 11.0	\$ 11.8
+ Vehicles	\$ 2.8	\$ 1.7	\$ 1.8
- Liabilities	\$ (12.3)	\$ (13.2)	\$ (13.7)
= Net Worth	\$ 104.8	\$ 80.4	\$ 87.1

Sources: Board of Governors of the Federal Reserve System, Survey of Consumer Finances (SCF) and Financial Accounts of the United States (FA). The SCF field period runs from the beginning of survey year Q2 through the end of survey year+1 Q1. Detailed reconciliation of SCF and FA balance sheet concepts is available from the authors.

A closer look at the divergence between SCF and FA balance sheet categories for 1995 and 2016 suggests more a nuanced explanation and an alternative approach to synthesizing the data (table 1). Again, the SCF is conducted over the course of a twelve-month field period, so we compare aggregates to both beginning and ending quarterly FA values. In the balance sheet categories where market prices are either easily observed or not relevant (financial assets and liabilities) the totals line up quite well at both the beginning and end of our sample.¹⁹ The

¹⁷ The SCF field period generally runs four quarters starting in the second quarter of the survey year. The connected squares line segments show the entire SCF field period, and helps add perspective about how much the FA values being compared can change while the SCF is in the field.

¹⁸ See, for example, Saez and Zucman (2016), Maki and Palumbo (2000), Sabelhaus and Pence (1999), and Cynamon and Fazzari (2016).

¹⁹ Some of the residual difference in liabilities, for example, is attributable to how certain types of debt are captured in the SCF. In particular, the SCF is missing some student debt for individuals outside the sample frame (living in

divergence between SCF and FA balance sheet aggregates is more pronounced in the three tangible asset categories where aggregate market values are not easily measured using available administrative or market data. Indeed, in 2016, the roughly \$20 trillion divergence between the SCF and FA net worth is almost entirely accounted for by real estate (the SCF finds about \$6 trillion more) and non-corporate business (the SCF finds about \$10 trillion more). Although quantitatively less important, the SCF also finds higher values for owned vehicles of about \$1 trillion.

It may seem obvious that the FA embodies the truth against which to benchmark the survey totals. Indeed, for financial assets and liabilities for which the FA aggregates are derived from source data from financial institutions, we deem that the FA is the appropriate benchmark. As such, we align the aggregate SCF level of financial assets and liabilities with the FA by scaling the individual amounts in the SCF accordingly. The decision to benchmark to the FA is less obvious in the asset categories with difficult to observe market values. In the case of owned real estate, for example, the FA is currently in the process of changing the methodology used to value those assets, and that change will eliminate much of the gap between FA and SCF housing values, raising the FA to be much closer to the SCF.²⁰ The gap between SCF and FA aggregates for equity in non-corporate businesses is attributable to a combination of conceptual and measurement differences, but those are not easily disentangled. The FA constructs the balance sheets of non-corporate businesses on a category-by-category basis, assigning market values to some assets such as real estate, for which price indexes exist.²¹ Other assets such as equipment and intangible property are valued at current cost. The net result of the conceptual and methodological differences is a much higher level of non-corporate equity in the SCF. Finally, the method used by FA to value the stock of owned vehicles involves multiplying price indexes by real stocks estimated using perpetual inventory methods, and either input could be problematic. In the SCF, car values are assigned from published NADA reports on a vehicle-by-vehicle basis. For all of these reasons, we do not benchmark SCF aggregates of housing, non-

student housing) and some of the household debt (in an FA accounting framework) of individuals running owned businesses. There are also likely unresolved issues with revolving credit, insofar as the source data for the FA is from financial institutions that do not distinguish convenience use of credit cards from true revolving debt outstanding.

²⁰ See Gallin et al. (2018). The fact that FA housing values were benchmarked to household survey reports prior to the early 2000s explains why the SCF and FA real estate numbers in Table 1 match quite well in 1995.

²¹ FA Table B.104 shows the balance sheet decomposition for the non-corporate business sector.

corporate business, and vehicles to FA values. However, we do use the FA aggregate saving measure for those components, effectively assuming that the differences in within-category wealth changes over time between the SCF and FA are due to differences in capital gains. From an agent-type and lifecycle perspective, benchmarking housing and vehicles to the FA in all periods would reduce wealth in the middle of the age and wealth distribution for whom housing is most important. Conversely, benchmarking equity in non-corporate businesses to the FA would dramatically lower wealth at the top of the wealth distribution.²²

Incomes in the NIPA and SCF

The combination of SCF micro and FA macro wealth data along with a method for backing out capital gains is sufficient to disaggregate saving using the left hand side of the intertemporal budget constraint ($S = \Delta W - G$). However, there are two reasons to incorporate micro and macro income data as well. First, we want to use the micro-level incomes to solve for consumption across birth cohorts and agent types, and consumption is the difference between income and saving ($C = Y - S$). Second, we need measures of income in order to create saving rates (S/Y) across age and agent-type groups. The steps needed to synthesize SCF micro and NIPA macro incomes are somewhat more involved than the steps for synthesizing balance sheets, because the NIPA aggregates include many imputed components not available in the SCF, and the SCF only asks about incomes in the year prior to each triennial survey.

We refer to the income concept that we seek to align as “adjusted disposable income.” The measure is effectively NIPA disposable income minus imputations for owner occupied housing, employer and government provided health insurance, and other in-kind transfers.²³ SCF and NIPA incomes and taxes are allocated nine categories, and the estimated aggregates for the first (1995-1998) and last (2013-2016) three-year periods in our sample are shown in table 2.²⁴ In total, the SCF captures roughly 80 percent of the corresponding NIPA incomes, but there is substantial variation across the income categories. In order to preserve the underlying distribution of each income component across birth cohort and agent type group, each SCF income component is scaled to match the NIPA independently.

²² Bricker et al. (2016) directly assess how the decision to benchmark affects wealth concentration estimates.

²³ See online appendix 1 for details about the steps taken to align NIPA and SCF income concepts.

²⁴ In table 2, SCF income for the year prior to the survey is multiplied by three in order to approximate the total over the three-year period.

Table 2. Adjusted Disposable Income in the NIPA and SCF (Trillions)			
Subperiod 1995 Q2 through 1998 Q1			
Disposable Income Components	SCF	NIPA	Percent Ratio
Wages and Salaries	\$ 12,051	\$ 11,278	107%
+ Business Income	\$ 2,175	\$ 2,370	92%
+ Social Security	\$ 675	\$ 1,047	64%
+ Interest and Dividends	\$ 712	\$ 3,216	22%
+ Other Government Cash Transfers	\$ 194	\$ 583	33%
+ Employer Retirement Contributions	\$ 272	\$ 833	33%
+ Retirement Interest and Dividends	\$ -	\$ 1,064	NA
- Personal Income Taxes	\$ 2,731	\$ 3,529	77%
- Employer Payroll Taxes	\$ 855	\$ 1,003	85%
= Adjusted Disposable Income	\$ 12,493	\$ 15,859	79%
Subperiod 2013 Q2 through 2016 Q1			
Disposable Income Components	SCF	NIPA	Percent Ratio
Wages and Salaries	\$ 24,133	\$ 22,906	105%
+ Business Income	\$ 5,125	\$ 4,839	106%
+ Social Security	\$ 2,250	\$ 2,554	88%
+ Interest and Dividends	\$ 1,276	\$ 5,220	24%
+ Other Government Cash Transfers	\$ 518	\$ 1,485	35%
+ Employer Retirement Contributions	\$ 579	\$ 1,496	39%
+ Retirement Interest and Dividends	\$ -	\$ 1,849	NA
- Personal Income Taxes	\$ 5,595	\$ 5,530	101%
- Employer Payroll Taxes	\$ 1,394	\$ 1,669	84%
= Adjusted Disposable Income	\$ 26,893	\$ 33,150	81%

Sources: Board of Governors of the Federal Reserve System, Survey of Consumer Finances (SCF), and Bureau of Economic Analysis, National Income and Product Accounts (NIPA). The SCF field period runs from the beginning of survey year Q2 through the end of survey year+1 Q1. Detailed reconciliation of SCF and NIPA concepts is available from the authors.

The largest components of income (like wages and business income) and taxes (estimated based on reported incomes) are well captured in the survey, and relatively little scaling is required. The two components with the largest gaps are at the bottom (other government cash transfers) and top (interest and dividends) of the income and wealth distribution. In both cases the SCF only captures about one-fourth to one-third of the corresponding NIPA value. The under-reporting of cash transfers in surveys is a well-known phenomenon, and difficulties capturing interest and dividends is likely related to the fact that most survey respondents see those flows as simply being rolled over, and not “income” per se. In any event, scaling the

observed SCF incomes to match the corresponding NIPA total is biased for our purposes only if there is differential reporting (relative to truth) across age and agent type groups.

Measuring retirement income flows in an internally consistent way is important for our disaggregation. Like most surveys, the SCF asks respondents about the incomes they receive from retirement plans, including both traditional pension plan benefits and withdrawals from account-type pensions. However, the measure of income that is consistent with our intertemporal budget constraint is the new contributions to retirement plans along with the interest and dividends earned on those plans. Note that on the left-hand side of the budget constraint we are measuring the saving in retirement plans as the change in retirement plan balances less capital gains. Although perhaps counterintuitive, the benefits paid and withdrawals from retirement accounts are dissaving, not income.

The SCF includes questions about employee and employer contributions to retirement plans. The employee contributions are subtractions from wages and business income, so (assuming respondents report incomes before those deductions as the survey requests) those contributions are captured as part of the underlying incomes. The employer contributions are not included in the usual SCF (or any other survey) income measures, but there are questions about such contributions in the labor force modules. The SCF captures only about 40 percent of employer contributions (table 2) because respondents are generally not knowledgeable about how much their employers are actually contributing to (especially traditional pension) retirement accounts. As with the other under-reported incomes, the key to our benchmarking strategy is that there is no differential reporting of employer contributions to retirement plans across age or agent-type groups.

Lastly, the SCF has no data on the interest and dividends earned on retirement accounts. Similar to the issue with reported dividends and interest mentioned above, most SCF respondents have little if any knowledge about how much their retirement account earns. Respondents do have a good sense of the balances in those accounts (as described in the previous section). Unlike taxable dividends and interest, however, the SCF does not even attempt to capture that information (table 2). Therefore, we allocate the missing interest and dividends using the reported retirement account balances.

Accounting for Interfamily Transfers

Interfamily transfers net to zero for the household sector as a whole, and thus play no role in NIPA or FA intertemporal budget accounting. However, interfamily transfers in any given year are the same order of magnitude as total household sector saving.²⁵ Interfamily transfers vary systematically by age and are highly unequal across the agent type groups in our disaggregation, and thus have differential impacts on the intertemporal budget constraints across the different agent types at different points in the lifecycle.²⁶ The important question when introducing transfers into the disaggregated intertemporal budget constraint identities is whether such transfers are measured well in the SCF data. That question is difficult to answer because, unlike income and wealth measures, there are no aggregate benchmarks.

There are two principal forms of interfamily transfers. The first and largest form of interfamily transfer is inheritances at death. The second form of transfers is *inter vivos* gifts and support. The *inter vivos* gifts and support can be further subdivided into alimony and child support versus voluntary transfers. In addition to the different forms of interfamily transfers, each form has both a giver and a receiver, and accounting for both flows is important in order to rearrange the disaggregated (birth cohort and agent-type) intertemporal budget constraints and thus disaggregate saving and consumption. Looking at transfers from both the giver and receiver perspectives also provides a data check in terms of internal consistency.

The SCF survey instrument directly captures three of the four transfer flows required for the intertemporal budget constraint disaggregation.²⁷ The SCF asks respondents about inheritances received, gifts and support paid, and gifts and support received, alimony and child support paid, and alimony and child support received. The missing element in the interfamily transfers identities is bequests made, which we estimate using a model of differential mortality and adjustments for inheritance taxes, funeral expenses, and other death-related costs.²⁸ Also, the bequest made by a deceased individual does not have a one-to-one correspondence with reported

²⁵ Feiveson and Sabelhaus (2018).

²⁶ The focus here is on direct transfers because we are disaggregating wealth change, but other indirect forms of wealth transmission are also certainly important. The SCF does contain questions that shed light on some of these channels, such as investment in education and inclusion in lucrative family businesses. See Feiveson and Sabelhaus (2018) for a discussion of how important these indirect channels are likely to be for explaining intergenerational wealth correlations.

²⁷ Online Appendix 2 provides details about the interfamily transfer measures described in this section.

²⁸ Consistent with the notation introduced in the next section, the SCF does not ask about inheritances received by surviving spouses.

inheritances received in the SCF, because any given decedent often has more than one heir. Therefore, we divide bequests by the number of living children in order to simulate what we expect to find in terms of reported inheritances.

	Percent of Total			
	Bequests Made		Inheritances Received	
	Count	Dollars	Count	Dollars
<50K	49	5	55	6
50K-299K	36	25	30	21
300K-599K	8	17	8	17
600K-1M	4	17	4	16
>1M	3	36	2	40
	(Thousands)	(Billions of \$)	(Thousands)	(Billions of \$)
Annualized Average	2,030	340	1,733	\$ 287

Source: Author's calculations using Survey of Consumer Finances (SCF) and other sources, see online Appendix 2 for details.

Table 3 compares the estimated net bequests with reported SCF inheritances over the 1996 through 2016 period. The results show that reported SCF inheritances received align well with our estimated bequests made, both in aggregate and across several size buckets. We estimate that on average, 2 million bequests are made each year, while 1.7 million inheritances are reported. The total dollars that flow across families are estimated at \$340 billion per year from the bequest side, and \$240 billion per year from the inheritances side. The remaining divergence is either attributable to under reported inheritances or misspecification in the mortality or other assumptions used to generate bequests, but in any case, the survey is internally consistent in terms of capturing transfers at death.

The distributions of transfers at death by size from the giver (bequest) and receiver (inheritance) perspectives line up quite well. Approximately half of all estimated bequests and reported inheritances are for amounts below \$50,000, but those account for only 5 or 6 percent of the dollars transferred. At the other end of the size distribution, the 6 or 7 percent of transfers above \$600,000 account for about half of all dollars transferred. This skewness interacts with our cohort and agent-type disaggregation in a predictable way, as Feiveson and Sabelhaus (2018) show that probability of receiving an inheritance and the size of that inheritance are both strongly correlated with lifecycle position (age) and the characteristics (like permanent income and education) that we use to define our agent-type groups below.

	Percent of Total			
	Support Paid		Support Received	
	Count	Dollars	Count	Dollars
<50K	91	54	93	62
50K-299K	8	26	6	25
300K-599K	0	5	1	9
600K-1M	0	7	0	4
>1M	0	8	0	0
	(Thousands)	(Billions of \$)	(Thousands)	(Billions of \$)
Annualized Average	5,726	\$ 54	5,865	\$ 43

Source: Author's calculations using Survey of Consumer Finances (SCF) and other sources, see online Appendix 2 for details.

The SCF survey instrument captures both sides of *intervivos* transfers directly. Table 4 shows alimony and child support transfers by size over the sample period. In total, there are roughly 5.7 million households reporting having paid alimony and child support in an average year, and 5.9 million reporting having received alimony and child support. The distributions by size also line up quite well, except for some large reported payments (probably one-time settlements) in the support paid category. Alimony and child support payments generally have only a second-order effect on our estimated saving and consumption profiles, because most of those transfers are within a birth cohort and agent-type group.

	Percent of Total			
	Gift or Support Paid		Gift or Support Received	
	Count	Dollars	Count	Dollars
<50K	85	38	73	10
50K-299K	13	34	20	21
300K-599K	1	13	4	13
600K-1M	0	6	1	8
>1M	0	10	2	49
	(Thousands)	(Billions of \$)	(Thousands)	(Billions of \$)
Annualized Average	12,688	\$ 155	442	\$ 48

Source: Author's calculations using Survey of Consumer Finances (SCF) and other sources, see online Appendix 2 for details.

The third category of interfamily transfers is *intervivos* gifts and support other than alimony and child support. Table 5 shows the distribution of these voluntary transfers across the

sample period, and in this case there is a clear conceptual divergence. The SCF captures gifts and support received in two modules. The first is the inheritance module, which is also the basis for the entries in Table 3 above. Respondents are asked about any “substantial” transfers received, and whether the transfer was an inheritance or a gift. The second point in the survey where gifts and support received are captured is in the income module.

The SCF *intervivos* transfers made information is collected after the income module, right after the questions about alimony and child support paid. The question asks if the respondent “provided any substantial financial support” to others, with an interviewer note to “include substantial gifts.” Clearly, the amount and number of gifts and support paid are both higher than corresponding amounts received, because the transfers made is a much broader concept. There is evidence that the large *intervivos* transfers of greatest interest are captured well from both perspectives. In terms of transfers made, the SCF finds about \$25 billion (16 percent of \$155 billion) above \$600,000. On the transfers received side, the SCF finds \$27 billion (57 percent of \$48 billion) above \$600,000. The divergence in the lower transfer size categories reflects at least one important conceptual difference, because respondents making transfers likely include non-cash transfers (such as college tuition or rent paid for someone else), while respondents receiving those same transfers do not report those as transfers received, because they are generally not cash or asset transfers. In the empirical work we calibrate the voluntary *intervivos* transfers to match the mounts received, which means we are capturing the high value transfers of interest, while effectively treating smaller (especially in-kind) transfers as the consumption of the individual making the transfer.

3. Pseudo-Panel Methodology

The SCF provides a series of representative and comprehensive snapshots of U.S. household balance sheets every three years. In this section we explain our methodology for disaggregating saving and consumption across groups (agent types and birth cohorts) and time. Agent “type” is kept intentionally vague at this point, but individual characteristics that do not change over time (such as educational attainment) or move slowly over time (such as permanent income) are the sorts of “types” the reader should initially have in mind. Relative to other types of research using pseudo-panel analysis, the biggest complications arise when measuring saving are because of (1) wealth transfers *between* groups, and (2) we only observe wealth holdings and incomes of individuals in the SCF if they are either the head of household or spouse/partner of the head of household.²⁹

The explanation of our methodology begins with what we observe in the SCF micro data, and how that relates to what we are trying to estimate. For each individual (head or spouse) we observe their net worth at time t , which we denote w_{it} . Although we will ultimately divide SCF net worth into several categories of wealth for assigning capital gains, we suppress the wealth type superscript and look only at total net worth to keep the notation simpler at this point. Most components of net worth in the SCF are reported as jointly owned when a spouse/partner is present, so we divide those equally. Incomes, transfers, and taxes are also divided equally. We also observe a vector of characteristics for every head and spouse, including the type of agent (j), their birth cohort (c), their marital/partner status ($m_{it} = 1$ if spouse/partner present, 0 otherwise), and the values of agent type and cohort for their spouse (j_s, c_s) if they have a spouse. We will also use other demographic and economic variables (x_{it}) that vary *within* agent type and cohort and affect differential mortality and the receipt of inheritances.

Timing

The goal of the exercise is to estimate savings and consumption across agent types using the balance sheet identity—namely, that the change in wealth can be decomposed between savings, gains, and net transfers to that group. (In the aggregate, the transfers net to zero.) To back out

²⁹ The SCF survey unit is a household, but detailed data is only collected on the Primary Economic Unit (PEU). Persons living in the unit who are reported as not financially interdependent, including roommates and adult children, are in the Non Primary Economic Unit (NPEU). The SCF collects only limited and highly aggregated data on individuals in the NPEU.

savings from this identity, we need to use the SCF to get estimates of the size and allocation of transfers (including bequests) and capital gains. Both of these depend on the timing of when the assets are acquired, and, as such, it is important to lay out the assumptions about timing we make when describing the pseudo-panel saving estimates:

- At the beginning of each three-year subperiod some individuals die, with a probability that depends on their agent type, cohort, and their own idiosyncratic characteristics associated with differential mortality within their type and cohort group.
- Non-mortality related entry and exit into an agent type and cohort group between t and $t+3$ also occurs at the beginning of the period. For example, children will move out of their parent's home, and become the head or spouse in a new household that is observed in the next survey wave.³⁰ We assume they bring zero wealth into the group total when they become a head or spouse at the beginning of the period, and we want to count their saving *during* the period and thus include them in the denominator when measuring average saving (along with average disposable income, transfers, and consumption). Also, older people may exit from head or spouse status if they (say) move in with their children. We assume that if they had any wealth, it is bequeathed at that point, meaning their wealth effectively gets the same treatment as if they died.
- Lastly, and consistent with the timing of deaths and entry/exit, all wealth is transferred (bequests made and received, as well as *intervivos* gifts made and received) at the beginning of the three year period. This implies that the capital gains that will accrue on that transferred wealth during the three year period will be credited to the group receiving the bequest at the beginning of the current three-year period. The key assumption when separating the change in wealth into the saving and capital gains components is that the capital gains *rate* is fixed by asset type within each time period.

Bequests Made

The distribution of bequests made is part deterministic, and part estimated. The deterministic part is associated with spousal bequests, because we know the agent type and

³⁰ The SCF has very little information about income and wealth on household members other than the head and spouse. In addition to children, non-surveyed roommates will also transition to head or spouse. In the SCF, only one roommate in a household will be in the PEU.

cohort of every married individual's spouse. The estimated part is due to bequests from single individuals who die. The bequests made by single individuals are put into a bequest "pool" from non-spouses which is allocated across all potential heirs, using an inheritance function that captures both the probability of receiving an inheritance and the amount received. As with the mortality function, the inheritance function has some features that are common to all members of agent type j and cohort c , at time t , but there could also be heterogeneity within agent types.

Denote every individual's probability of death between time t and time $t+3$ using $d(j,c,t,x_{it})$, where j =agent type, c =cohort, t = year, and x_{it} is the vector of individual characteristics that affect differential mortality. Then, the total amount of bequests at death *made* by agents of type j in cohort c , at the beginning of the time $(t, t+3)$ sub-period is given by,

$$B_{jc(t,t+3)}^- = \sum_{i \in j,c} w_{it} d(j,c,t,x_{it})$$

Total bequests made by all individuals ($B_{(t,t+3)}^-$) because of death is just the sum over all agent and cohort types, which is,

$$B_{(t,t+3)}^- = \sum_j \sum_c B_{jc(t,t+3)}^-$$

Since our final goal is to measure the savings across a 3-year period of the survivors—i.e. those individuals who did not die in that time period—it is useful to define the wealth of the survivors in group j,c in time t . Denoting $W_{jct} = \sum_{i \in j,c} w_{it}$:

$$W_{jct}^{\text{survivors}} = W_{jct} - B_{(t,t+3)}^-$$

Net Transfers Received

The next step is to determine the amount of bequests *received* by the surviving individuals in each cohort and agent type group. Those bequests that accrue to agent type js and cohort cs at time t through the direct spousal link is:

$$B_{js,cs,(t,t+3)}^{+sp} = \sum_i w_{it} d(j,c,t,x_{it}) m_{it} (1 - d^s(js,cs,t,x_{it}^s))$$

Where js and cs are the observed agent type and cohort of the *spouse* of individual i , and $d^s(js, cs, t, x_{it}^s)$ is the mortality function for the spouse (the spouse has to be alive in order to receive the bequest). The total pool of non-spousal bequests is given by,

$$B_{(t,t+3)}^{+ns} = B_{(t,t+3)}^{-} - \sum_j \sum_c B_{jc(t,t+3)}^{+sp}$$

These remaining bequests are distributed across all other surviving individuals.³¹

Non-spousal bequests are allocated across agent type and cohort groups using inheritance functions, $b^{+ns}(j,c,t,x_{it})$, which, like the mortality functions, have both group-level and individual-specific inputs. These functions are derived from the self-reports of inheritances received in the SCF. The mortality-adjusted inheritance function of individual i is $b^{+ns}(j,c,t,x_{it}) * (1 - d(j,c,t,x_{it}))$. The condition on inheritances received is simply that the sum across all individuals equals the pool of non-spousal bequests, $B_{(t,t+3)}^{+ns}$. That is,

$$B_{(t,t+3)}^{+ns} = \sum_{j,c} \sum_i b^{+ns}(j, c, t, x_{it}) (1 - d(j, c, t, x_{it}))$$

The amount of non-spousal bequests received by the j, c group at time t (B_{jct}^{+ns}) is then just the sum of these calibrated amounts for individuals of agent type j in cohort c .

Similarly, we defined V_{jct}^{+} as the *inter vivos* transfers received by the j, c group and V_{jct}^{-} to be the *inter vivos* transfers given by the j, c group. As with bequests received, these are calculated from the self-reports of respondents as described in the appendix. Thus, the total net interfamily transfers received by the survivors of the j, c group at time t is:

$$IFT_{jct(t,t+3)} = B_{jct(t,t+3)}^{+sp} + B_{jct(t,t+3)}^{+ns} + V_{jct(t,t+3)}^{+} - V_{jct(t,t+3)}^{-}$$

Capital Gains

The last step when working with the change in total wealth for the groups comprised of agent types j and cohorts c is to determine the capital gains ($G_{jct(t,t+3)}$) accruing to each group. At this point, we (trivially) expand our notation to include asset and liability categories, adding a superscript z to each wealth variable when we intend to break it down into different categories. If we assume that each asset and liability category has its own capital gains rate ($g_{(t,t+3)}^z$) estimated (as described above) using aggregate data on asset prices, then the total capital gains earned by the survivors in group j, c between time t and $t+3$ is,

$$G_{jct(t,t+3)} = \sum_z (W_{jct}^{survivors,z} + IFT_{jct(t,t+3)}^z) g_{(t,t+3)}^z$$

³¹ The distributed bequests are allocated across age and agent-type groups as described in Feiveson and Sabelhaus (2018). Inheritances received are adjusted for estate taxes and other costs. Online appendix 2 has details about how those adjustments and other aspects of the interfamily transfer accounting are implemented in practice.

Saving and Consumption

Finally, with all pieces in place, we can now use the intertemporal budget identity to back out saving and consumption for each group:

$$S_{jc(t,t+3)} = W_{jct+3} - W_{jct}^{\text{survivors}} - \text{IFT}_{jc(t,t+3)} - G_{jc(t,t+3)}$$

The saving of survivors is the difference between their wealth in time $t+3$ and their wealth in time t minus their other sources of wealth flows, namely their net transfers and their capital gains. Note that if we aggregate this identity across all groups, total *inter vivos* transfers given and received offset, as do the bequests given (which is subtracted from the survivors' wealth) and the inheritances received, leaving the aggregate identity in the macro data. That is, aggregate savings is equal to the change in wealth minus total capital gains. Thus, group-level saving rates estimated using the above equation will add up to the aggregates familiar to macroeconomists. (Since we assume that death occurs at the beginning of the 3-year period, the saving of non-survivors is zero, which is why the saving of survivors sums to the aggregate.)

Having solved for saving, if we add the micro data on total disposable income ($Y_{jc(t,t+3)}$) for cohort c and agent type j over the three-year sub-period windows, the other side of the intertemporal budget constraint can be used to solve for consumption:

$$C_{jc(t,t+3)} = Y_{jc(t,t+3)} - S_{jc(t,t+3)}$$

These equations for saving and consumption are precisely the decomposition we will take to the pseudo-panel for the empirical results presented in the next two sections. Every component of the saving and consumption disaggregation *except* capital gains is determined completely by the synthesized micro and macro data. In addition to solving for the levels of saving and consumption across cohort and agent-type groups, the question remains how best to exhibit and describe the outcomes of interest. One approach is to compute and report per-capita values for the various measures across cohort groups, agent types, and sub-periods, which we do in the next section. In addition, we also construct two measures of saving rates across cohorts, agent types, and time. The first is $(S_{jc(t,t+3)}/Y_{jc(t,t+3)})$, which is a direct decomposition of the usual NIPA and FA saving concepts discussed in the previous section. The second saving measure is $(S_{jc(t,t+3)} + \text{IFT}_{jc(t,t+3)} + G_{jc(t,t+3)})/(Y_{jc(t,t+3)} + \text{IFT}_{jc(t,t+3)} + G_{jc(t,t+3)})$, which captures the flows of total resources not consumed.

4. Sources of Wealth Change over the Lifecycle

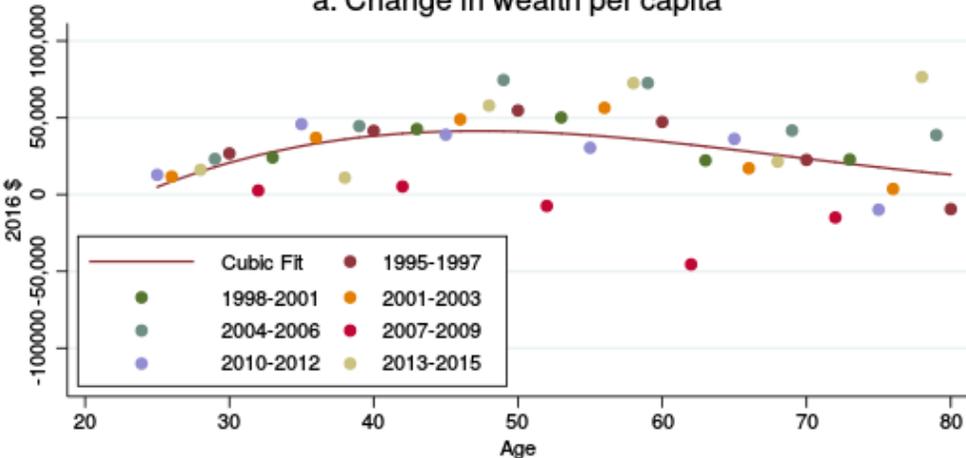
The pseudo-panel decomposition presented in the previous section provides the empirical framework one needs to construct estimates for lifecycle patterns of saving and wealth accumulation, but there is one statistical obstacle. The SCF is a relatively small sample, and although the sampling strategy is careful to represent families at the top of the wealth distribution, there is no stratification by age within wealth groups. Thus, inferences about sources of wealth change across birth cohorts and three-year sub-periods is subject to extensive sampling variability. In this section, we group SCF observations by 10-year birth cohorts, and show that despite the sampling variability, there is clear evidence of the sort of textbook lifecycle patterns that are difficult to find in other data sets and empirical approaches.

The first set of results (figure 4) are simple scatterplots of the change in wealth and the sources of change in wealth and other flows for birth cohorts born between 1910-19 and 1980-89. Each birth cohort is represented as a point estimate in between one and seven of the three-year SCF sub-periods, because the oldest cohorts are not included at the end of the sample period (they are mostly deceased by 2013-2015) and the youngest cohorts are not included at the beginning (they were not in their own households in 1995-1997). In total, there are 40 data points in panel of figure 4, each representing a unique birth cohort and three-year sub-period. Every ten-year birth cohort/three-year sub-period point estimate is plotted along the horizontal (age) axis at the midpoint of their age range at the time of the survey. For example, the midpoint age of 1960-69 birth cohort was 30 in 1995, and it was 51 in 2016.

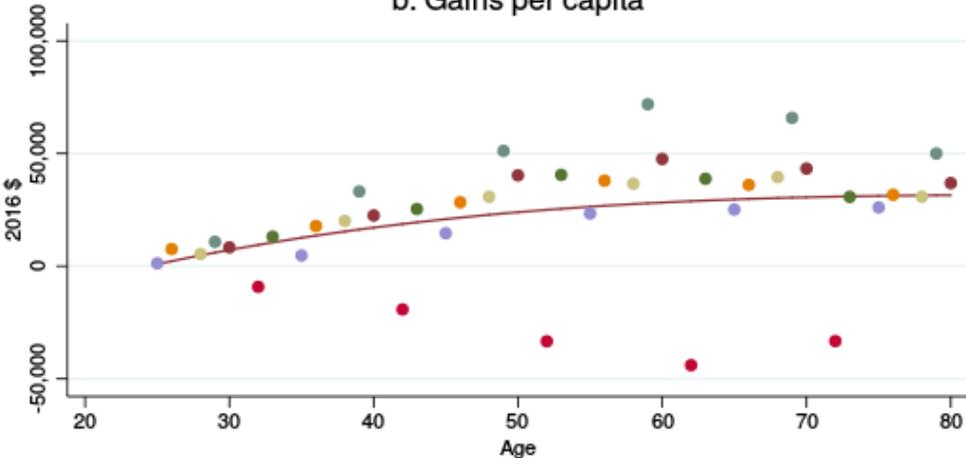
The line on each graph in Figure 4 shows the smoothed lifecycle profile of that component when we pool over all cohorts and across all time periods. That the dots cluster around the smoothed line for some components—namely transfers, income, saving, and consumption—suggests that the lifecycle shape of these components does not vary greatly over the business cycle. However, there are some stark outliers for the wealth changes that are driven by cross-time divergences in capital gains. The color-coding of the dots, which distinguishes time periods, helps to show that the three-year sub-period 2007-2009 stands out in terms of widespread capital losses, while 2004-2006 stands out in terms of strong capital gains. Our empirical strategy assures that total gains will sum to the aggregate in any three-year sub-period, and the scatterplot shows that gains and losses are highly correlated across sub-groups.

Figure 4. Per capita flows

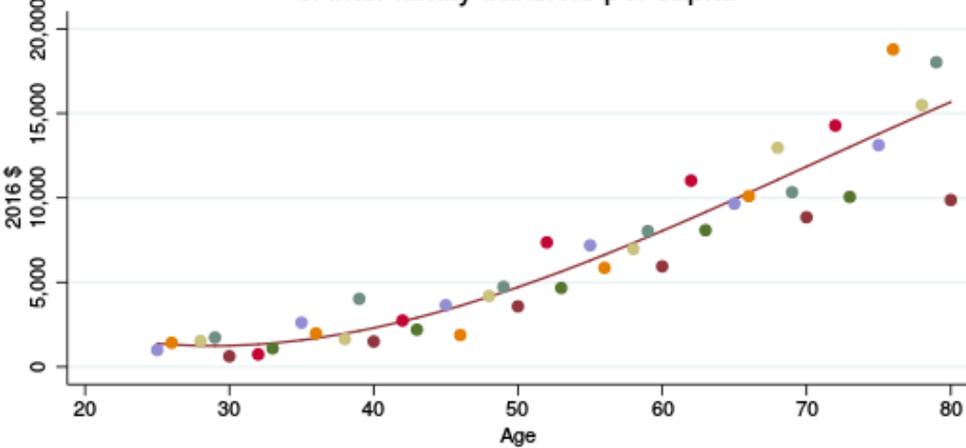
a. Change In wealth per capita

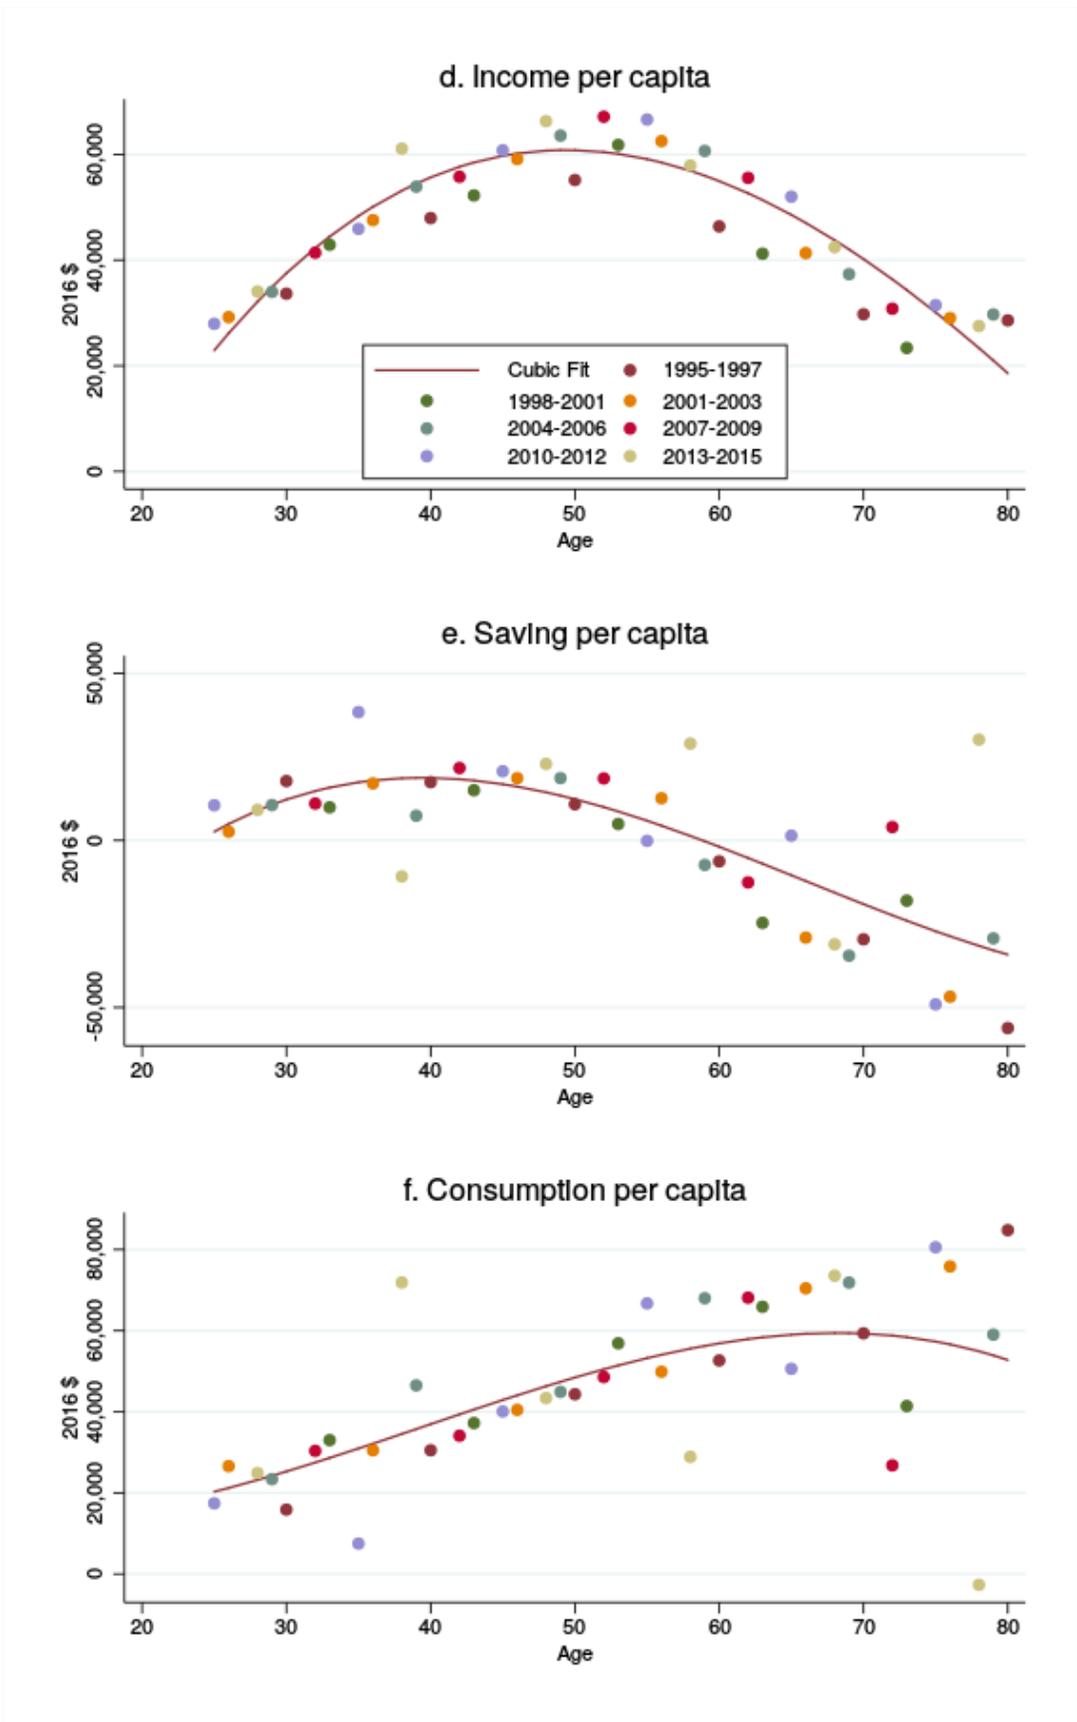


b. Gains per capita



c. Inter-family transfers per capita





Despite the sampling variability in the SCF data, there is clear evidence of the sorts of lifecycle patterns generally shown in economics textbooks, though with some important real-world caveats. Per-capita wealth (panel a) is clearly increasing most rapidly through middle age, but, unlike the textbook model, wealth is still increasing (the wealth change is >0) at the end of the lifecycle. The only exception to the positive wealth change is three-year sub-periods where capital losses dominated the other sources of wealth change. The generally positive wealth accumulation at all ages is generally consistent with findings from research using other data and approaches, where the failure to spend down wealth at older ages is often interpreted as evidence against lifecycle behavior.

The decomposition of wealth change into component flows using the intertemporal budget constraint in figure 4 helps make it clear why data and theory seem at odds. The lifecycle pattern of per-capita disposable income (panel d) has the usual hump-shape shown in textbooks, despite the bias from looking at multiple cohorts on the same chart.³² The other key ingredient in the textbook charts is negative saving at older ages, which shows up clearly in panel e. Average per-capita saving turns negative just before age 60, and becomes increasingly more negative over the remainder of the lifecycle.

A simple textbook model of lifecycle saving and wealth accumulation would be challenged by these conflicting findings. How can saving turn negative (as the lifecycle model suggests) while wealth continues to grow? The mechanical answer lies in panels b and c, per-capita capital gains and interfamily transfers. Per-capita capital gains are (on average) positive throughout the lifecycle, and interfamily transfers received are steeply increasing with age. Most inheritances are received at older ages, but more importantly, the death of one spouse has the immediate effect of doubling the per-capita wealth of the surviving spouse. The sum of gains and interfamily transfers is sufficient to offset the negative (conventionally measured) saving such that per-capita wealth continues to grow.

The deeper answer about why measured saving and wealth change seem disconnected involves reminding ourselves about the specific concepts of income, saving, and wealth shown in figure 4. The wealth concept in our decomposition is total market wealth owned by the household at a point in time, and that includes claims to future private pension benefits and other

³² All of the per-capita values in figure 4 are in real terms, but younger cohorts have higher lifetime earnings, and thus the true income trajectories for a given cohort are steeper and decline less at older ages.

retirement accounts. The income concept in our synthesized data is consistent with that wealth concept, because income includes the additions to retirement wealth in a given year, which is new contributions and the interest and dividends earned on those accounts. Benefit payments and withdrawals from retirement accounts are not, however, a source of income in the conceptually correct decomposition. A retiree may think of a pension benefit or 401(k) withdrawal as income, but it is actually a drawdown of accumulated wealth. That simple difference between these conceptually correct measures of income and saving and the concepts generally used in micro data helps explain why most research concludes that retirees do not dissave at older ages. The fact that retirees do not spend their entire pension benefits just means they are dissaving a little more slowly, not that they are saving.

The fact that positive average wealth change over the entire lifecycle is accounted for by capital gains and interfamily transfers focuses our attention on the concept of saving itself, and ties the findings here back to the results from registry data sets such as Bach, Calvet, and Sodini (2018) and Fagerang, Holm, Moll, and Natvik (2018). Both of those papers struggle with the exact same problem raised here. Should the increases in wealth over the lifecycle associated with capital gains and interfamily transfers be included in measured saving, in the sense of intended additions to wealth? The topics we ultimately care about are wealth concentration and comovements in income, consumption, and wealth. Our estimates of consumption, solved for using disposable income minus saving and shown in panel f of figure 4, indicate a steady increase in spending over the lifecycle. Individuals are clearly increasing their consumption as income and wealth accumulate over the lifecycle. Individuals are not, however, increasing consumption fast enough to spend down the increases in wealth, but some of that is not difficult to understand. For example, an older person whose house goes up in value may not be likely to spend against that wealth increase, because they want to continue living in the home and enjoying the flow of housing services, without taking on risk by using (say) a reverse mortgage. A similar argument applies that a business owner whose labor income is tied to the business asset might not consume more when that asset goes up in value, lest they forego the labor income.³³

³³ One other consideration worth noting is that the lifecycle wealth decomposition shown in figure 4 is completely data driven *except* for the allocation of capital gains across cohort and agent-type groups. To the extent that allocating gains proportionally by asset type across age groups is inappropriate, the misallocation flows through the intertemporal budget constraint to the saving measures, and thus (given income) to the consumption measures. An important component of the next steps for this project is testing whether the estimated lifecycle patterns are sensitive to the proportional capital gains assumption.

5. Lifecycle Patterns Across Permanent Income Groups

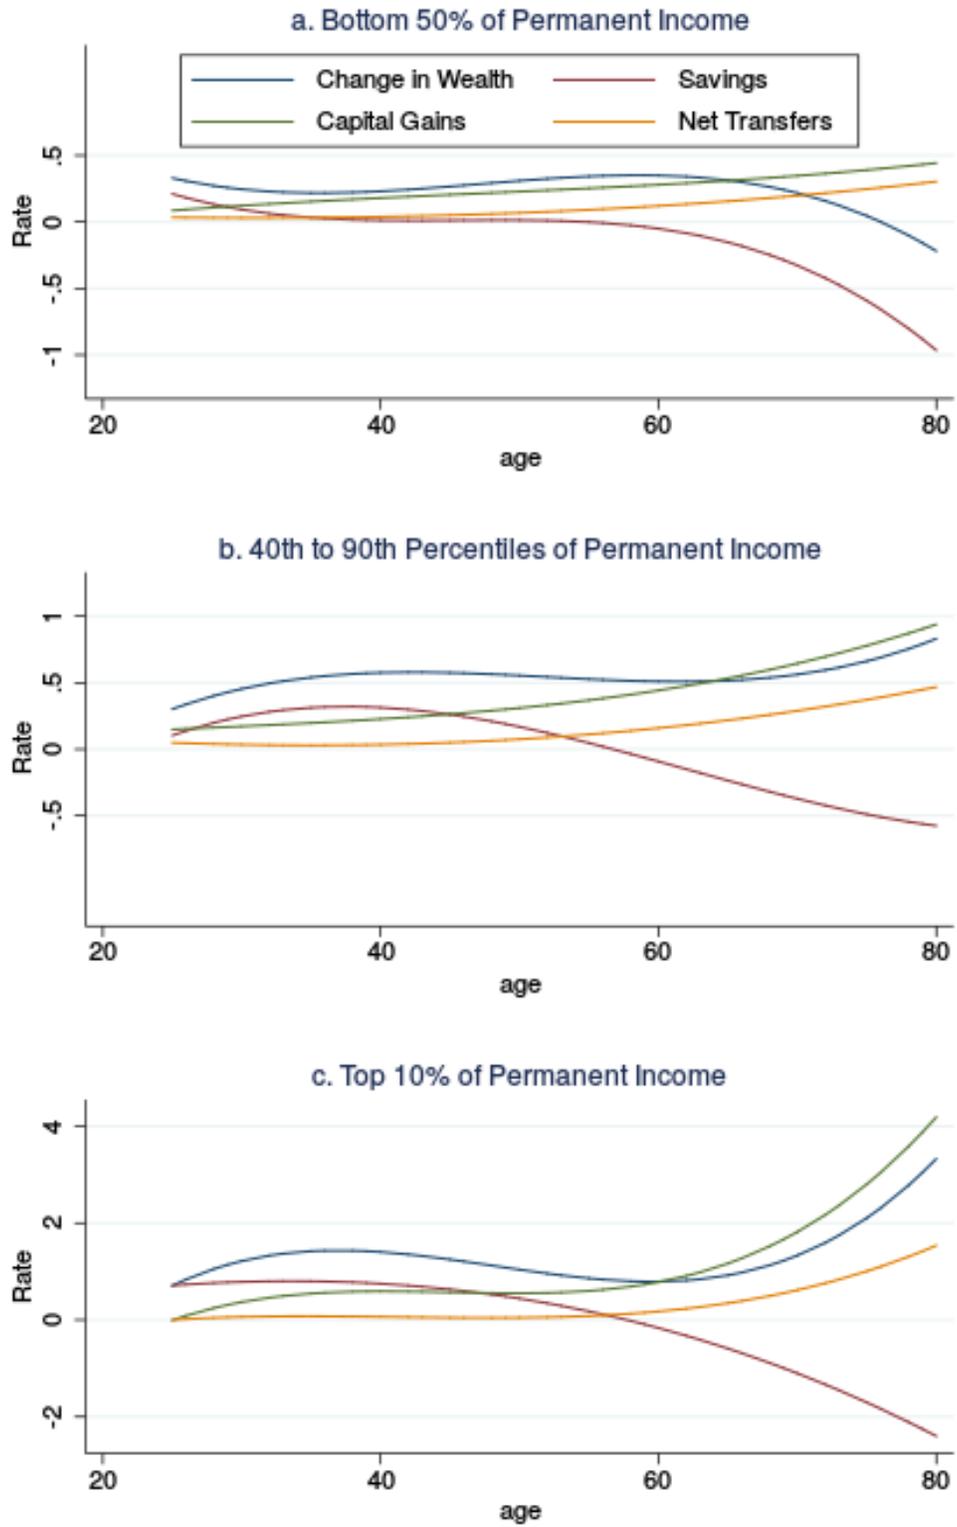
The conceptual and empirical framework developed above is applicable to disaggregating both across- and within-cohort sources of wealth change. Cohort-level wealth decomposition is systematic because age and time move together, but for any other identifiable agent-type group the movement of individuals across groups between the multiple cross-sections can be problematic for the pseudo-panel. In this section we use the SCF measure of permanent income to further disaggregate lifecycle saving and wealth accumulation.³⁴ We independently sort the members of each birth cohort by permanent income in each survey, and thus we link groups within cohort agent-types across time using the relative positions in the permanent income distribution.³⁵

Data limitations guide our choices about how to disaggregate cohorts into permanent income groups. The SCF sample sizes are between 4,500 and 6,500 families during the period we are studying, though the estimation sample is per-capita, so we effectively have more observations than the SCF sample size suggests because the SCF main respondent and spouse/partner enter separately. Even so, after splitting the sample by birth year, there is relatively little sample left to divide by permanent income within cohorts. However, the SCF oversampling strategy makes it possible to achieve a fair amount of statistical precision at the very top of the permanent income distribution. Thus, we divide each birth cohort into three permanent income groups: the bottom half of the population, the 50th to 90th percentiles, and the top 10 percent. SCF cross-sections show that the bottom half of the permanent income distribution holds few assets, and what they do own is mostly in the form of housing and vehicles largely by debts. The 50th through 90th percentiles of the permanent income distribution has modest wealth holding, especially as they accumulate retirement balances and pay down housing and vehicle debt. The top 10 percent both has incomes that far exceed the other groups, and they hold an even more disproportionate share of wealth, especially owned businesses.

³⁴ The SCF measure of permanent income separates transitory fluctuations by asking respondents if the reported income is what they would earn in a “normal” year. If the answer is no, they are asked what it would be in a normal year. Devlin-Foltz and Sabelhaus (2016) show that the statistical properties of the permanent and transitory decomposition are similar to those retrieved by backing out shocks using panel income data.

³⁵ Another agent-type to consider is education groups. The same decompositions below using education yields similar results. In on-going work we are also exploring wealth as agent-type, which is complicated because wealth is the outcome we are tracking across the cross-sections. However, it may be feasible to combine known distributions of income shocks and capital gains with an assumption about the functional relationship between income, consumption, and wealth to solve for the within wealth group saving rates without the extreme assumption (as in Saez and Zucman (2016)) that there are no movements across wealth groups between cross-sections.

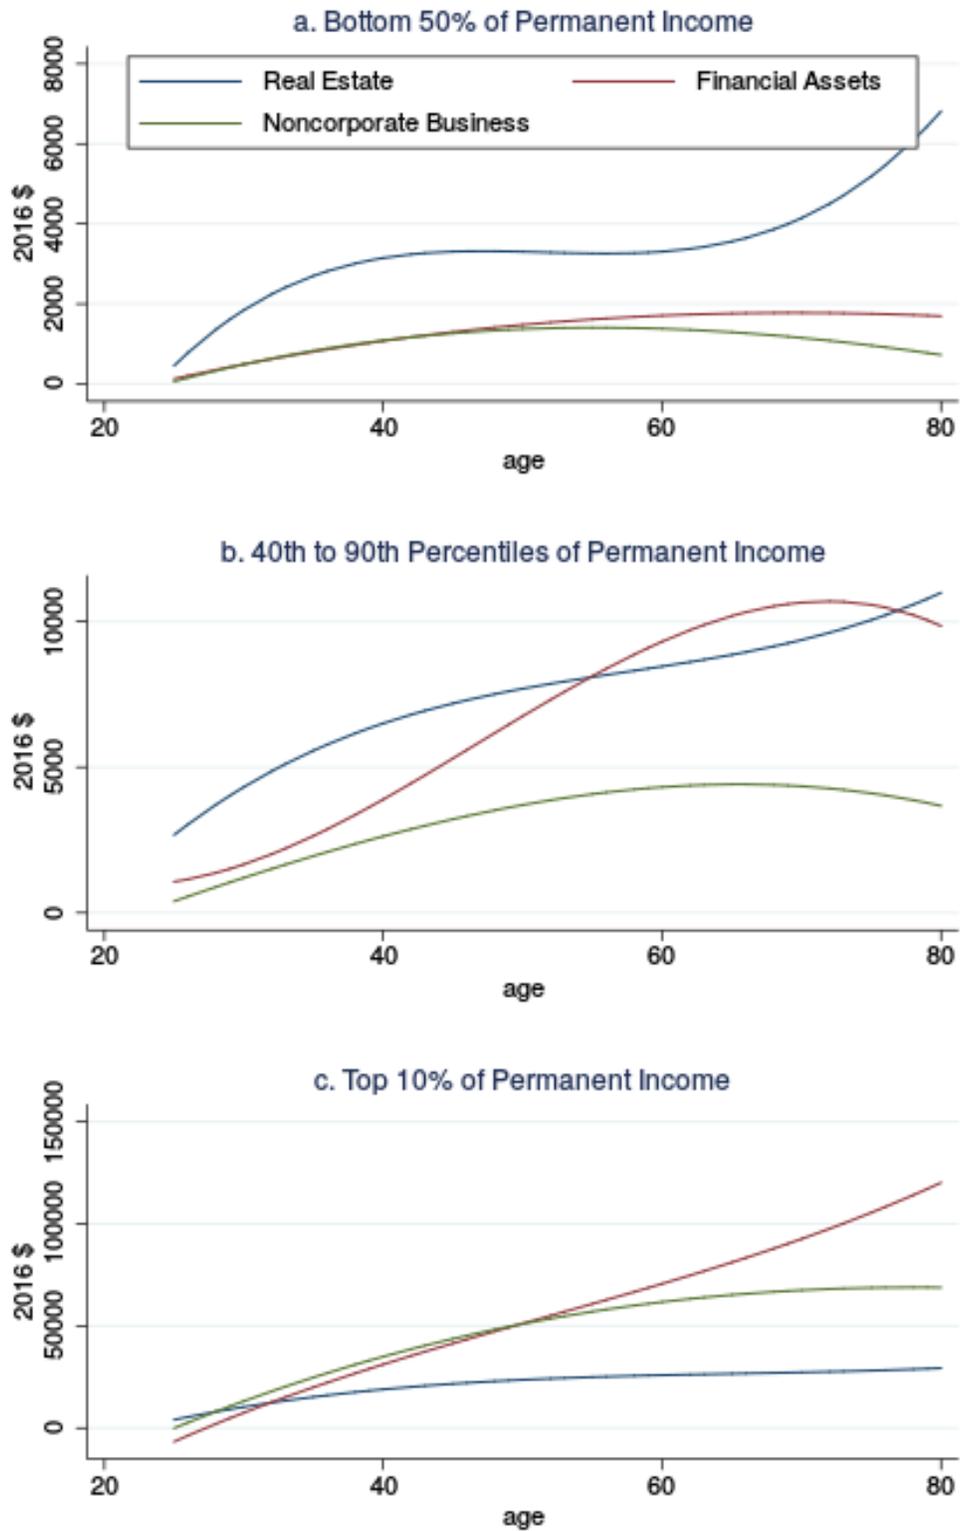
Figure 5. Wealth flows divided by income



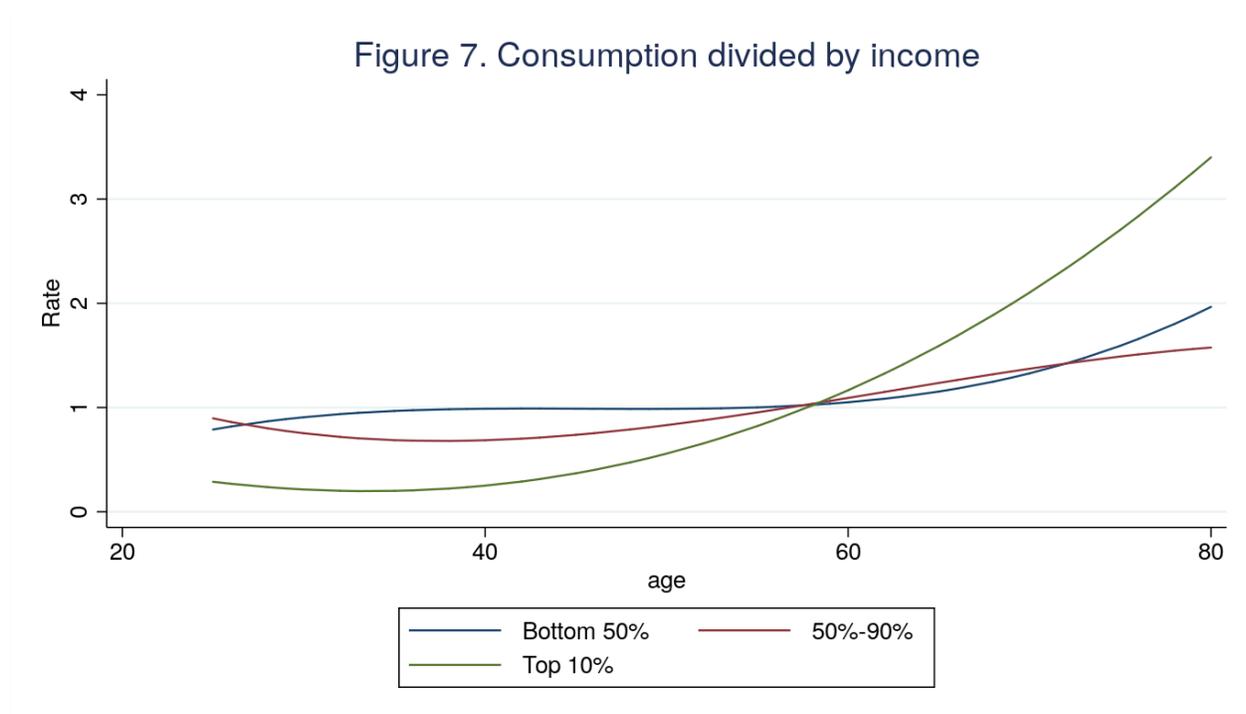
The first thing to note about the three panels of figure 5 is the scale. Even though the wealth change and components of wealth change in figure 5 are scaled by disposable income for each permanent income group, the vertical axis changes substantially as we move from the bottom half of the permanent income distribution to the top 10 percent. However, the second thing to note is that there are important similarities in lifecycle patterns across the three groups. In particular, saving turns negative around the same age for each group, while capital gains and net transfers received become increasingly important with age. The differences in curvature (and thus vertical scale) across the three income groups is driven by both changes in the numerator and the denominator with age. In the numerator, higher income individuals have more positive saving when young and more negative dissaving when old. In the denominators, higher income individuals see a larger relative drop in their disposable income as they get older, in part because a larger share of their retirement income is from pensions and other retirement account withdrawals, which are not part of disposable income. Lower income individuals are more reliant on Social Security, which is a component of disposable income.

The key takeaways about lifecycle saving and wealth accumulation in figure 5 come from looking at the *levels* of saving relative to income over the lifecycle. At the bottom of the permanent income distribution, saving hovers around zero for most age ranges, which is not surprising given the fact that wealth cross-sections show low net wealth holdings for low permanent-income groups. However, the bottom half of the income distribution does accumulate (and eventually spend down) some wealth, though much of that increase is in the form of capital gains. The 50th through 90th percentiles by permanent income have a more textbook-like saving rate pattern, with saving peaking around age 40. Finally, the top 10 percent of the permanent income distribution saves a much larger share of their incomes while young, and for a much longer period. Again, it is useful to remember that saving here includes forms of wealth accumulation that are generally missed in other survey-based saving measures, especially the build-up of retirement accounts and businesses that are not accounted for by capital gains or interfamily transfers received.

Figure 6. Gains per capita

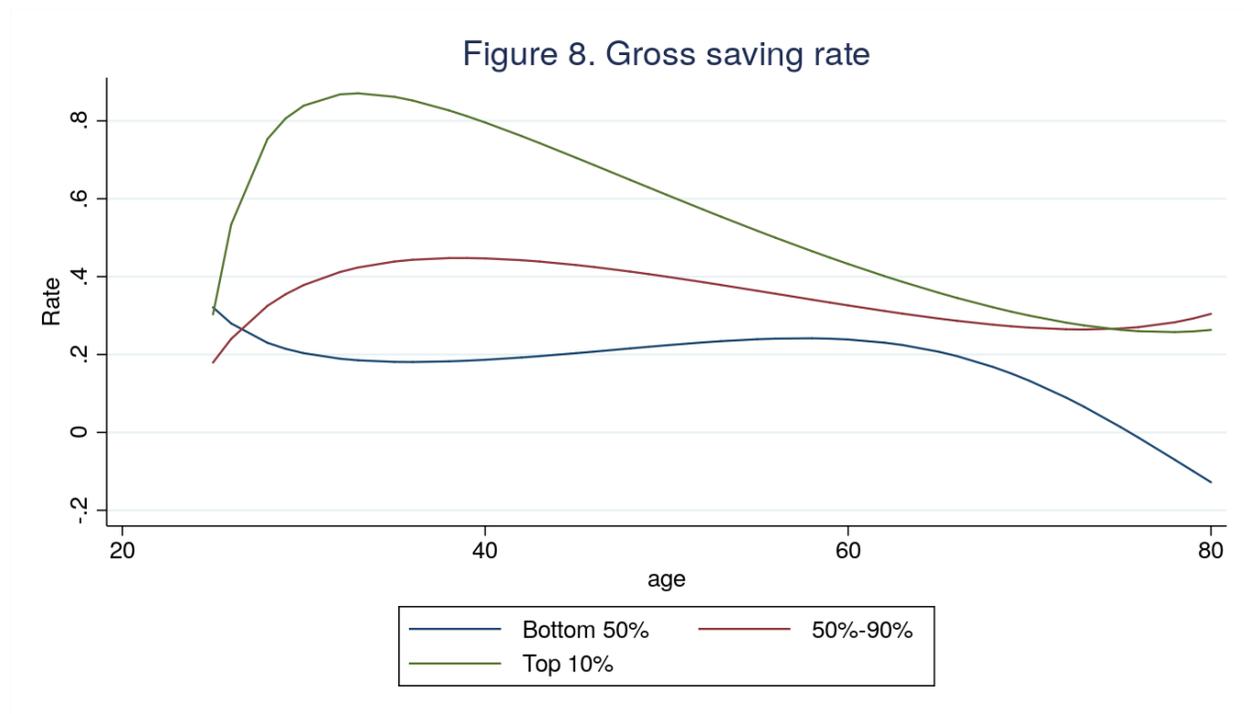


Although wealth levels differ dramatically across permanent income groups, one common feature is an increasing role for capital gains as a source of wealth change over the lifecycle. Figure 6 shows per-capita gains across the three asset classes for which we compute gains and across the three permanent income groups. Again, the vertical scale is very different, with the scale for the top 10 percent by permanent income and order of magnitude larger. The sources of capital gains are consistent with what we know about asset holdings, because real estate dominates capital gains in the bottom half, and is on par with financial assets in the 50th to 90th percentile. Capital gains for the top 10 percent of the permanent income distribution are dominated by financial assets and (through middle age) by owned businesses.



Returning to the right-hand side of the intertemporal budget constraint, the permanent income decomposition is applied to solve for consumption relative to income, which is effectively one minus the saving rates shown in figure 5. The three consumption/income ratios by age are shown in figure 6. The charts reinforces the messages of figure 5, meaning consumption is roughly equal to disposable income for the bottom half of the income distribution over working years, and slightly below income for the 50th to 90th percentiles. The shape of the lifecycle consumption/income profile for the top 10 percent by permanent income is extreme, and although the same considerations apply here (disposable income falls more for the top

income group) we cannot rule out that the method for distributing gains (proportional by asset holdings) may be affecting the pattern. In particular, if we are assigning too little gains for younger agents, that means (because wealth change is data) we are assigning too much saving, which means we are deriving too little consumption. Thus, the issue of proportionally distributing capital gains is a key question for the next steps in this research.³⁶



The final observations from the pseudo-panel decomposition by cohort and permanent income focus on the broader concept of the “gross” saving rate, mentioned in section 3. The gross saving rate is effectively one minus the share of total resource flows (disposable income, capital gains, and interfamily transfers) that is not consumed. Figure 8 shows that the relative patterns of saving rates across permanent income groups are in line with the NIPA and FA concept saving rates, but shifted up, and (for all but the bottom half income group) never turn negative. Again, this begs the question of whether assets like housing or businesses are part of wealth that lifecycle agents intend to draw down, but the empirical evidence shows that the gross and net concepts of saving are very different for all income groups at lifecycle frequencies.

³⁶ One other feature of the SCF worth noting is that there are several measures of consumption collected in the survey, as well as realized capital gains. It may be possible to use that information to help calibrate gains versus consumption across the age distribution.

6. Conclusions

This paper presents the first set of results from a new approach to disaggregating the change in wealth over the lifecycle for U.S. households. The approach is a pseudo-panel methodology, in which series of cross-section snapshots are synthesized with aggregate time-series to measure the contributions of saving, capital gains, and interfamily transfers to overall wealth change by age and for various agent types. The wealth change decompositions here are largely consistent with the lessons from individual level administrative panel data in countries where such data exists. In particular, the disaggregation here emphasizes the role of capital gains in accounting for wealth change over most of the life cycle, with savings playing a large positive role for some agent-types early in life, and a substantial negative role later.

The specific estimates presented here are based on one particular way to construct the pseudo-panel, fitting functional representations to measured changes for overlapping ten-year birth cohorts. However, the method can (and will be) generalized to more flexible specifications, such as kernel-smoothed profiles. In addition, the lifecycle estimates here are for the entire period spanned by our data (1995 through 2016) but the pseudo-panels can be estimated for each of the seven three-year sub-periods within that overall span as well. In that sense, the pseudo-panel approach can be used to disaggregate wealth change (and consumption and saving) over the lifecycle and business cycle frequencies for the various agent types.

In addition to refining and extending the pseudo-panel methodology, the steps required to synthesize the micro and macro data focus attention on conceptual aspects of lifecycle decision-making and measuring economic well-being that will guide this project going forward. The narrow focus on marketable wealth and the corresponding treatment of various income sources suggests that any distributional analysis based on those concepts will be missing the impact of important policies at lifecycle frequencies.³⁷ There is a direct connection between these sorts of measurement issues and saving rates: low and middle-income families do most of their lifecycle saving through social insurance, so measuring the change in net present value of social insurance on an accrual basis could fundamentally change inferences about overall saving behavior.³⁸

³⁷ The current state of market-based distributional analysis is well captured in the DINA approach of Piketty et al. (2018). Their cross-section estimates show that economic well-being for the young and middle-income families who are facing higher social insurance taxes is decreasing over time, while economic well-being for the older cohorts is rising. That divergence would not be observed in a comprehensive lifecycle measure.

³⁸ Devlin-Foltz, Henriques, and Sabelhaus (2016) show that counting future Social Security benefits in household wealth has a first-order equalizing impact on the distribution of retirement resources.

7. References

- Attanasio, Orazio, and Luigi Pistaferri. 2014. "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure," *American Economic Review*, 104(5): 122-126. (May)
- Bach, Laurent, Laurent E. Calvert, and Paolo Sodini. 2018. "From Saving Comes Having? Disentangling the Impact of Savings on Inequality." Working paper.
- Baker, Scott, Lorenz Kueng, Steffen Meyer, and Michaela Page. 2018. "Measurement Error in Imputed Consumption." Working paper.
- Benhabib, J., A. Bisin, and S. Zhu. 2015. "The Wealth Distribution in Bewley Economies with Capital Income Risk," *Journal of Economic Theory*, 159, 489–515.
- Benhabib, J., A. Bisin, and S. Zhu. 2017. "Earnings Inequality and Other Determinants of Wealth Inequality." *American Economic Review*, 107(5): 593-597. (May)
- Benhabib, J. and A. Bisin. 2018. "Skewed Wealth Distributions: Theory and Empirics," *Journal of Economic Literature*, 56(4):1261-1291. (December)
- Bosworth, Barry, Gary Burtless, and John Sabelhaus. 1991. "The Decline in Saving: Some Microeconomic Evidence," *Brookings Papers on Economic Activity*, 1:1991, p. 183-256.
- 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.
- Bricker, Jesse, Lisa J. Dettling, Alice Henriques, Joanne W. Hsu, Lindsay Jacobs, Kevin B. Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A. Windle. 2017. "Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, 103(3): 1-42. (September)
- Bricker, Jesse, Alice Henriques, and Peter Hansen. 2018. "How Much has Wealth Concentration Grown in the United States? A Re-Examination of Data from 2001-2013," Finance and Economics Discussion Series 2018-024. Board of Governors of the Federal Reserve System (U.S.).
- Carroll, Christopher D., Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. "The Distribution of Wealth and the Marginal Propensity to Consume." *Quantitative Economics*, 8(3): 977-1020.
- Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull. 2003. "Accounting for the U.S. Earnings and Wealth Inequality." *Journal of Political Economy*, 111 (August): 818–57.
- Cynamon, Barry, and Steve Fazzari. 2016. "Household Income, Demand, and Saving: Deriving Macro Data with Micro Data Concepts." *Review of Income and Wealth*, 63(1): 53-69.

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." Discussion Paper No 11746, Center for Economic Policy Research.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz Pardo. 2016. "The Implications of Richer Earnings Dynamics for Consumption and Wealth." Working Paper No. 21917, Cambridge, MA: National Bureau of Economic Research.

Detting, Lisa J., Sebastian Devlin-Foltz, Jacob Krimmel, Sarah Pack, and Jeff Thompson. 2015. "Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances," Federal Reserve Board: FEDS Working Paper 2015-86. (October)

Devlin-Foltz, Sebastian, and John Sabelhaus. 2016. "Heterogeneity in Economic Shocks and Household Spending in the US," *Fiscal Studies*, 37(1): 153-192 (March).

Devlin-Foltz, Sebastian, Alice Henriques, and John Sabelhaus. 2016. "Is the US Retirement System Contributing to Rising Wealth Inequality?" *Russell Sage Foundation Journal of the Social Sciences*, 2(6) 59-85.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy*, 112(2): 397-444.

Fagereng, Andreas, Luigi Guiso, Davido Malacrino, and Luigi Pistaferri. 2016. "Heterogeneity and Persistence in Returns to Wealth." Working Paper No. 22822, Cambridge, MA: National Bureau of Economic Research.

Fagerang, Andreas, Martin Holm, Benjamin Moll, and Gisle Natvik. 2018. "Saving Behavior across the Wealth Distribution." Presentation, NBER Summer Institute 2018.

Feiveson, Laura, and John Sabelhaus. 2018. "How Does Intergenerational Wealth Transmission Affect Wealth Concentration?," FEDS Notes. Washington: Board of Governors of the Federal Reserve System, June 1, 2018, <https://doi.org/10.17016/2380-7172.2209..>

Fisher, Jonathan, David Johnson, Jonathan P. Latner, Timothy Smeeding, and Jeffrey Thompson. 2016a. "Inequality and Mobility Using Income, Consumption, and Wealth for the Same Individuals." *Russell Sage Foundation Journal of the Social Sciences*, 2(6):44-58.

Fisher, Jonathan, David Johnson and Timothy Smeeding. 2016b. "Inequality of Income and Consumption in the U.S.: Measuring the Trends in Inequality from 1984 to 2011 for the Same Individuals." *The Review of Income and Wealth*, 61(4), 630 – 50.

Gallin, Joshua H., Raven Molloy, Eric Nielsen, Paul Smith, and Kamila Sommer. 2018. "Measuring Aggregate Housing Wealth: New Insights from an Automated Valuation Model," Finance and Economics Discussion Series 2018-064. Board of Governors of the Federal Reserve System (U.S.).

- Gabaix, Xavier Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica*, 84(6): 2071–2111.
- Gale, William G., and John Sabelhaus. 1999. "Perspectives on the Household Saving Rate," *Brookings Papers on Economic Activity*, 1:1999, p. 181-224.
- Kopczuk, Wojciech. 2015. "What Do We Know about the Evolution of Top Wealth Shares in the United States?" *Journal of Economic Perspectives* 29(1): 47–66.
- Krusell, Per and Anthony A. Smith, Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy," *The Journal of Political Economy*, 106(5):867-896.
- Maki, Dean M. and Michael G. Palumbo. 2001. "Disentangling the Wealth Effect: A Cohort Analysis of Household Saving in the 1990s," *Federal Reserve Board Finance and Economics Discussion Paper Series No 2001-21*.
- Mian, Atif, and Amir Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 101(5): 2132-56. (August).
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018 "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics* 133: 553-609. (May)
- Sabelhaus, John, and Karen Pence. 1999. "Household Saving in the '90s: Evidence From Cross-Sectional Wealth Surveys," *Review of Income and Wealth*, (45): 435-53. (December)
- 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.

Online Appendix 1. Synthesizing Micro and Macro Saving and Wealth Measures

Overview

This appendix describes how the Financial Accounts (FA), National Income and Product Accounts (NIPA), and Survey of Consumer Finances (SCF) data sources are used to construct the disaggregated saving and wealth accumulation measures in the paper. The three data sources are all essential for completing the intertemporal budget constraint in both the macro and micro data. However, the published concepts of income, wealth, and saving in each of the data sources are different from the conceptually adjusted measures shown in the paper. Here, we describe the conceptual adjustments and show how the adjusted macro estimates align with the adjusted and aggregated micro measures.³⁹

The integrated micro/macro framework starts with the intertemporal budget constraint, which states that the change in wealth (ΔW) is equal to the sum of saving (S) and capital gains (G). Given that saving is also the difference between disposable income (Y) and consumption (C), it is also true that $\Delta W = Y - C + G$. The table below illustrates how the information in each of the data sets aligns with the elements of the budget constraint identity in the paper.

Wealth, Income, Capital Gains, Consumption, and Saving in Various Data Sets

	Change in Wealth (ΔW)	Disposable Income (Y)	Capital Gains (G)	Saving (S)	Consumption (C)
Financial Accounts (FA)	ΔW^{FA}		G^{FA}	S^{FA}	
National Income and Product Accounts (NIPA)		Y^{NIPA}		S^{NIPA}	C^{NIPA}
Survey of Consumer Finances (SCF)	ΔW^{SCF} (some W^{SCF} components scaled to match W^{FA})	Y^{SCF} scaled to match Y^{NIPA}	G^{SCF} $= \Delta W^{SCF} - S^{SCF}$	S^{SCF} scaled to match S^{FA}	C^{SCF} $= Y^{SCF} - S^{SCF}$ $= Y^{NIPA} - S^{FA}$

The first thing to note is that FA and NIPA each have measures of aggregate saving (denoted S^{FA} and S^{NIPA}) that are in principle capturing the same thing. Indeed, the Federal Reserve Board regularly publishes the two saving measures, and they do track each other well over time, though

³⁹ The Stata code used to construct the SCF, FA, and NIPA measures with explicit references to FA/NIPA tables and line numbers and SCF variable names is available from the authors upon request.

the FA measures is somewhat noisier.⁴⁰ The second thing to note is that the SCF does not have an independent saving measure; some assumption about capital gains is required to solve for saving. The SCF wealth levels are used as a starting point, SCF saving is scaled to match the FA measure, and SCF capital gains are then solved for by subtracting FA saving from SCF wealth change (where some components of SCF wealth are scaled to match FA, more below).

In principle, because SCF saving is scaled to match FA total saving, we could then solve for SCF consumption by starting with the SCF income measures and subtracting the scaled saving measures. However, the first problem is that the SCF does not attempt to measure incomes over the entire three-year period between surveys. Other measurement and conceptual problems also suggest that starting with unadjusted SCF incomes when solving for disaggregated consumption is problematic. Thus, as indicated in the table, we scale the SCF incomes to match NIPA values. The derived SCF micro-level consumption values are thus a hybrid: aggregated micro consumption is the difference between NIPA disposable income and FA saving. The fact that FA and NIPA saving are very close in aggregate suggests that the hybrid micro-level consumption will also be close to NIPA consumption.

The notable difference between the disaggregated micro values and published aggregates is in capital gains. As shown in the text, most aggregated SCF balance sheet measures track well with the corresponding FA values, after conceptual adjustments. However, there are a few categories with substantial differences, most notably owned real estate, non-corporate business equity, and vehicles. In these cases, the SCF aggregates are higher and have grown faster than FA aggregates over time. One way to deal with this would be to scale the SCF to match the FA measures. If sampling or household-level measurement problems were the source of divergence, that might be the preferred approach (as it is for SCF versus NIPA incomes and other SCF balance sheet components, such as debts). In this case, though, there are good reasons to believe that the SCF measures are unbiased. Thus, by allowing SCF to determine the total change in wealth that is to be disaggregated and setting the levels of saving in each asset type to match the FA totals, we are allowing the aggregate capital gains measures in the SCF to diverge from the FA capital gains.

The discussion above about macroeconomic aggregates skirted over another set of issues also addressed below. How and why do the *concepts* of wealth, income, consumption, and saving differ across the various data sets? The first issue is coverage. The SCF sample covers households in a strict sense, but both NIPA and FA cover the household *sector*, which includes households per se and non-profit institutions serving households. The second issue involves the meaning of wealth, income, consumption, and saving. For example, certain balance sheet items in the FA are not collected in the SCF, and there is no good way to impute those measures. There are also in-kind consumption (such as government or employer provided health care) and some imputed incomes in the NIPA that (again) are not collected in the SCF. Some of these could in principle be imputed, but most involve offsetting income and consumption entries (what the health care was worth in terms of income matches what was spent on health care) and thus have no impact on the net saving measure we are focused on. Finally, in addition to aligning consumption and saving concepts in the SCF, FA, and NIPA, there are also additional issues that arise when *distributing* consumption and saving, because of retirement account flows and capital gains realizations.

⁴⁰ See FRB Z1 release, Table F.6, lines 41 and 45, available at <https://www.federalreserve.gov/releases/z1/current>.

Balance Sheets

The FA and SCF asset and liability data are well described elsewhere, and we will not repeat the careful conceptual and empirical reconciliation work that has been done in those other papers.⁴¹ Rather, we focus here on the specific adjustments made for each of the data sets, how the adjusted data compare to published values, how the adjusted SCF and FA track each other over time, and the implications of the adjustments for aggregate saving and capital gains.

Our starting point for the adjusted FA balance sheet measures is published Table B.101.h.⁴² Table B.101.h is a relatively new table (published beginning September 2018) that separates households from non-profit institutions. The historical and still published Table B.101 has households and non-profits together in the comprehensive household “sector.” The other key FA table used is R.101, which separates the change in several asset categories into saving (investment) versus capital gains (“revaluations” in FA parlance). Table R.101 has a third component of the first differences in asset balances, called “other volume changes.” The other volume changes are associated with changes in source data, disaster losses, and charge-offs. Given that our framework is focused on saving versus capital gains, we choose to lump (the usually small) other volume changes with the saving component, by subtracting revaluations from wealth change and labeling the residual saving. The gains measures in published Table R.101 are also adjusted to remove non-profits (there is no published Table R.101.h).

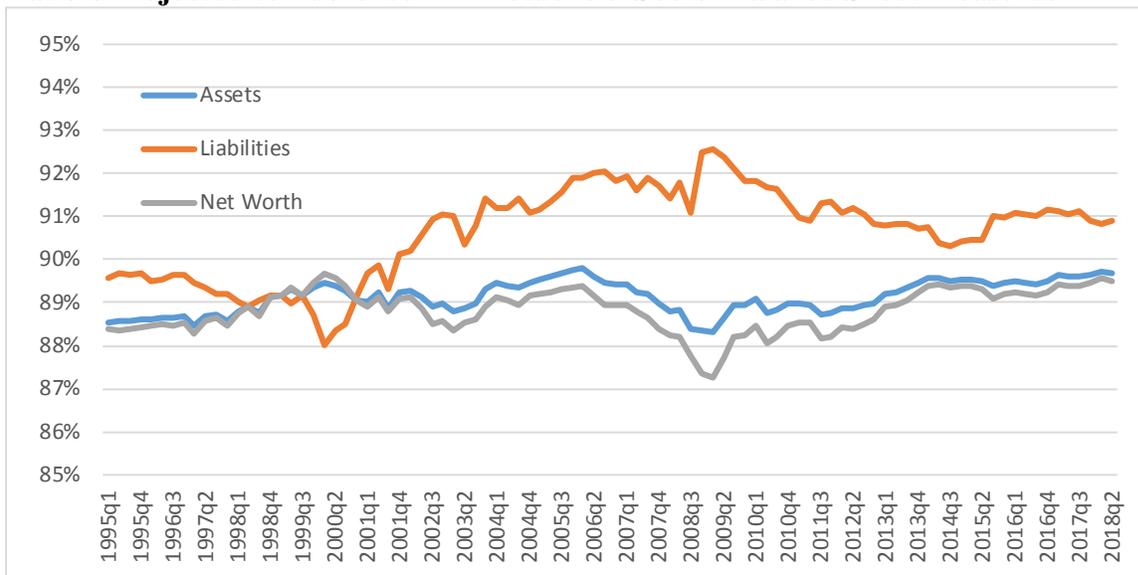
The FA balance sheet measures in Table B.101.h are very close to SCF concepts, but a few adjustments are needed to bring assets, liabilities, and net worth into conceptual alignment. First, there is no measure of the value of “other durables” in the SCF, so we exclude Table B.101.h line 6, and thus the consumer durables category is limited to vehicles. Second, miscellaneous assets category (line 23) includes things like the reserves that insurers hold for conditional payouts on health and other policies, and there is no corresponding entry in the SCF, so that is dropped from the FA as well. On the liabilities side, there are three main categories within “other loans and advances” (line 29) which are separated and treated differently. The largest piece is loans against (for example) 401(k)s and similar accounts. The SCF reports such accounts on the asset side net of loans, so we net those out against the asset side pension measure in the FA. The second largest piece is margin loans on brokerage accounts, which has a similar offset on the asset side. The residual in the “other” category is government agency loans to households, and those are captured in the SCF as liabilities, so no adjustment is needed. Finally, “deferred and unpaid life insurance premiums” (line 30) is not measured in the SCF. In principle, this liability could be netted out against “miscellaneous assets,” but (see above) those are already excluded from the asset side of the balance sheet. Dropping both implies we are not valuing household’s (contingent) claims on insurance companies or the premiums that households owe to life insurance companies.

⁴¹ See 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, and Dettling, Lisa J., Sebastian Devlin-Foltz, Jacob Krimmel, Sarah Pack, and Jeff Thompson. 2015. “Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances,” Federal Reserve Board: FEDS Working Paper 2015-86. (October).

⁴² All published FA series are accessed through the Federal Reserve Board web site at <https://www.federalreserve.gov/releases/z1/current>.

The adjustments to B.101.h concepts needed to match SCF concepts are, on net, relatively small. The most notable change is dropping other durables, which lowers household net worth by about three percent. The other adjustments net out to less than one percent for household net worth. Thus, overall, relative to published household *sector* net worth, the important adjustments involve removing the non-profit holdings from the balance sheet totals in B.101 which we do by starting with B.101.h. The ratios of our adjusted FA assets, liabilities, and net worth to B.101 published totals are shown below. Both the asset and net worth ratios are stable and near ninety percent. That is, taking out non-profits (and the other small adjustments, like subtracting consumer durables) lowers assets and net worth by about ten percent of the sector total in every period. The liability ratio is slight higher, around ninety-one percent in recent years.

Ratio of Adjusted to Published FA Household Sector Balance Sheet Measures



There are a few notable changes to the SCF balance sheet required to bring the micro concepts closer to the adjusted macro concepts. The most important adjustments include assigning a value for defined benefit (DB) pensions, adding the assets and liabilities of household members outside the primary economic unit, and adjusting consumer debt for timing and coverage differences.⁴³

⁴³ Details on the SCF DB pension imputation are available in Devlin-Foltz et al (2016). The non-primary economic unit (NPEU) income and wealth values are a growing share of the respective aggregates, and the code for allocating those NPEU measures is available from the authors. The SCF concept of debt diverges from the FA concept because the SCF focuses on debt balances “after the most recent payment” while the FA is for a point in time. Thus, the FA includes revolving debt that will be paid off (from an SCF perspective) within the billing cycle. Also, the SCF attempts to isolate household and owned-business credit, and that distinction is less bright in the aggregate data. The code that implements these specific adjustments within the SCF is available from the authors.

Disposable Income and Consumption

As with the balance sheet items in the FA, the published NIPA income, consumption, and saving for the household sector includes non-profit institutions serving households. In addition, the NIPA has a number of imputations for the value of goods and services such as owned housing and financial services, and offsetting income and consumption entries (especially for health insurance and other in-kind goods) that are difficult to measure at the household level. The fact that they have offsetting entries for income and consumption means that the level of saving is not affected.

The interconnection between non-profit institution and household income and consumption is somewhat more complicated than for balance sheets, as described in NIPA Table 2.9. There are substantial flows back and forth between the two sectors, most of which are irrelevant for computing saving and consumption (if a household gives donations to a non-profit that then spends the money, it is effectively consumption of the household). The only required adjustment is to subtract interest and dividends earned by non-profits (lines 50 and 51) from the totals to solve for capital incomes of households.

Housing leads to a substantial imputation in the NIPA, but the impact on saving is neutral. In the NIPA, owner occupied housing generates an imputed rental income flow that is the difference between what the owned housing stock would rent for on the open market and the costs (interest, maintenance, property taxes) that homeowners face. We exclude the imputed rent on owner occupied housing (NIPA Table 7.9) from both income and consumption. That implicitly (and appropriately) leaves the costs of maintaining owner-occupied as part of consumption, but leaves aggregate saving unaffected.

In-kind goods and services provided by employers and government (such as health insurance, housing assistance, and food stamps) also have offsetting effects on both income and consumption. We use NIPA Table 2.1 and 3.12 to identify employer insurance payments and government in-kind transfers, then use NIPA Table 7.20 to separate employer contributions to pensions from the total employer insurance payments in Table 2.1. A comprehensive measure of income or consumption would include these flows, but given our focus here on saving, subtracting the in-kind entries from both income and consumption appropriately leaves aggregate saving unaffected. The offsetting adjustments average about fifteen percent of income (or outlays) over the 1995 to 2016 period, and there is a downward trend (about six percentage points) as health and other in-kind transfers have grown in importance.

One apparent inconsistency between published aggregate and micro level income and consumption measures is only relevant for disaggregation, but worth mentioning here for completeness. The NIPA and FA aggregate saving concepts both treat retirement accounts in the same appropriate way. The change in retirement account balances is new contributions plus interest and dividends less benefit payments and withdrawals (the “saving” part) and capital gains. In the micro data, we observe the change in retirement account balances, and using an estimate of capital gains, we derive a measure of retirement account saving (or dissaving) at the micro level.

Online Appendix 2. Interfamily Transfers in the Survey of Consumer Finances

Overview

This appendix provides details about the data and methods used to estimate interfamily transfers. The goal is to quantify two different types of transfers (bequests/inheritances versus *inter vivos* transfers) from two different perspectives (who is receiving the transfer versus who is making the transfer). This two-by-two view of interfamily transfers makes it possible to check the internal consistency of the estimated flows, and to identify the demographic and economic characteristics of families making and receiving transfers. The primary data source is the triennial Survey of Consumer Finances (SCF) for the eight waves conducted between 1995 and 2016.⁴⁴ We also use published estate tax data from IRS, cohort mortality rates from the Social Security Administration (SSA) and estimates of demographic mortality differentials from the Congressional Budget Office (CBO). The overall strategy is summarized in the table below:

SCF Interfamily Transfers Estimation Strategy		
	Bequests/Inheritances	<i>Inter vivos</i> Gifts and Support
Transfers Received	<ul style="list-style-type: none"> • Reported inheritances received during previous three years from inheritance and gift module • Reported inheritance income in previous year from income module “other income” question code 12 • Supplemented with incremental information about real estate and business assets received in previous three years as inheritance or gifts (i.e. not purchased) 	<ul style="list-style-type: none"> • Reported alimony and child support received in previous year from income module support question • Reported other forms of support and gifts received in previous year from income module “other income” question codes 14, 28 • Reported gifts received during previous three years from inheritance and gift module
Transfers Made	<ul style="list-style-type: none"> • Estimated using survey year wealth holdings multiplied by demographically-adjusted (i.e. differential) three-year cumulative mortality rates • Bequests assigned to surviving spouse if present otherwise divided equally by number of children 	<ul style="list-style-type: none"> • Reported alimony and child support paid in previous year from income module follow-up • Reported other support paid and substantial gifts made in previous year from income module follow-up

The summary table above reveals a number of important aspects of the estimation strategy, each covered in more detail in the subsequent sections. A few preliminary observations are worth noting:

⁴⁴ For an introduction to the SCF and overview of 2016 results, see Jesse Bricker, Lisa J. Dettling, Alice Henriques, Joanne W. Hsu, Lindsay Jacobs, Kevin B. Moore, Sarah Pack, John Sabelhaus, Jeffrey Thompson, and Richard A. Windle. 2017. “Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances,” *Federal Reserve Bulletin*, 103(3). Available at www.federalreserve.gov/publications/files/scf17.pdf.

- The SCF directly measures three of the four cells in the two-by-two matrix, but the fourth (bequests made at death) has to be estimated because of the inherent difficulties with interviewing the deceased (or the representatives of their estates). The SCF concept of inheritances received does not include spousal transfers, so we distinguish bequests made at death by single people from those of married couples.
- There are key timing differences in the period over which various flows are measured. Inheritances and large gifts received can in principle be measured over a respondent's lifetime, because the questions are worded to cover inheritances and gifts *ever* received. In practice we use a three year look-back period for most of our inheritance and gift analysis, in order to span the time periods between surveys. However, most of the other flows in the two-by-two matrix are for the "previous" year (meaning survey year -1) the timing of which coincides with the SCF core income module questions.
- Various flows are intermingled and captured in different parts of the survey, and there is some redundancy that serves as back up for capturing potentially missed transfers. For example, *inter vivos* transfers received are captured in the inheritance module (if the respondent reports the amount as a substantial gift) and/or in the income module (as regular alimony or child support income or as a component of "other" income). Also, in the real estate and business modules, respondents are asked whether they purchased the asset or received it as a gift or inheritance. In principle, those transfers should be captured in the inheritance module as well, but in some cases there is incremental information in the asset section because the transfer is not captured in the inheritance module.
- Alimony and child support paid and received are both asked about separately in the SCF, so in principle the *inter vivos* column be separated between child support/alimony and other forms of *inter vivos* transfers. However, there is potential overlap with other forms of regular support that are captured using the SCF "other" income variable or the income module follow-up about support provided and "substantial" gifts to others.
- The only substantial conceptual gap between transfers made and received is in the *inter vivos* category. Reported "other support paid" and "substantial gifts made" is an order of magnitude larger than the corresponding reported *inter vivos* receipts, where measured receipts include lifetime substantial gifts received in the inheritance and gift module plus the amount of other support during the past year in the income module. The key to understanding this divergence is the failure of the survey to capture support received. For example, college students and their parents may view those "transfers" differently.

The remainder of this appendix provides details about how estimated bequests are measured, how the SSA baseline mortality rates and the CBO differential mortality model were applied to the SCF cross-sections, how the relevant SCF modules were used to construct our estimated transfer flows, and the specific wording of SCF questions.

Estimating Bequests Made at Death

Bequests made at death from the giver's perspective are not captured in the SCF survey, so we estimate bequests using SCF wealth holdings, cohort mortality rates from the Social Security Administration (SSA), and mortality differentials estimated by the Congressional Budget Office

(CBO) for their long-term microsimulation (CBOLT) model. The bequest estimates are generated in a way that conceptually matches what is being observed on the inheritance side, which makes it possible to check the internal consistency of transfers at death from the perspective of transfers made and received. In particular, one can look at the number of transfers and dollars transferred by size of bequest made and inheritance received to see if they line up.

The most important conceptual adjustment involves who makes bequests. The SCF “inheritances received” module instructs respondents to explicitly rule out transfers received from a deceased spouse, so we differentiate between bequests of single people and those of married couples. This distinction is also consistent with the IRS estate tax data we use to adjust bequests, because most non-taxable estates claim the one-time spouse deduction, and thus we use only the data from taxable estates to estimate deductions (more below). In short, bequests are only (probabilistically) generated if a single person dies or both members of the couple die in the same three-year period. Otherwise, in a married couple, the bequest if one member dies (meaning half of the family’s net worth) is designated to be a transfer to the surviving spouse. The fraction of estimated transfers at death going to spouses is generally about 10 to 25 percent higher than the estimated amount going to bequests in every three-year time period. That is, just under half of wealth-weighted deaths generate bequests, and the rest generate spousal transfers.

Adjusted Bequests

The starting point for estimating the level of bequests made is SCF net worth. The concept of wealth we use to estimate bequests begins with the SCF “Bulletin” net worth measure, which does not count non-transferable wealth such as the present value of defined-benefit pensions.⁴⁵ In addition to directly transferable assets, we also add the face value of life insurance to the potential estate.

There are three adjustments that drive a wedge between potential bequests made and inheritances received, particularly for wealthy decedents. All estates face some basic costs, such as funerals and expenses associated with distributing real assets. In addition, high end estates often make large charitable contributions, pay very high fees for executors and lawyers, and pay estate taxes. We use data on funeral expenses and other costs to adjust expected bequests for the vast majority of SCF cases, and we use data from the IRS Statistics of Income (SOI) to more carefully adjust predicted bequests for high end estates.

The definition of high end estates for our purposes is largely driven by estate tax rules and the associated data published by SOI.⁴⁶ Estate tax rules have varied over time, but the data are published in a consistent way back to 1995. In 2016, for example, the estate tax filing threshold was \$5 million, having risen from \$600,000 in 2003 and earlier. According to SOI, there were about 12,000 estates that filed in 2016, but of those, only about 5,000 were taxable. Much of the gap between taxable and non-taxable estate counts is accounted for by spousal deductions, because estates where one member of the couple dies generally choose to pay no tax by using the

⁴⁵ The “Bulletin” net worth concept is so named because that definition is the one published as the top line wealth number in the triennial SCF publications in the *Federal Reserve Bulletin*. The SCF *Bulletin* concepts are defined in the SAS macro available at www.federalreserve.gov/econres/files/bulletin.macro.txt.

⁴⁶ Available at <https://www.irs.gov/statistics/soi-tax-stats-estate-tax-statistics>.

spousal deduction. In that sense, the SOI taxable estates data is closest to our estimates, because we only generate expected bequests for single people and married couples that both die in the same year. The 5,000 taxable estates above \$5 million accounted for something like \$108 billion in wealth transfers, but of that, \$23 billion was spousal transfers (some couples choose to pay some of the tax when only one member of the couple dies, to avoid the progressive rate structure) so the relevant benchmark for gross estates above \$5 million in 2016 is something like \$85 billion. Our corresponding estimate of gross estates in the SCF is a bit higher, but that is somewhat expected, because assets (especially business holdings with no observable market price) are generally valued below market for estate tax purposes.

We use published SOI estate tax data to generate four adjustments applied to SCF predicted bequests above the filing threshold in each year. The published SOI data has gross estate size classes that vary by year. For example, in 2016 there are five size classes ranging from less than \$5 million to \$50 million or more. In the earlier years there are as many as eight size classes, ranging from less than \$1 million to \$20 million or more. For each gross estate size class in each year, we compute (1) the fraction of estates that claim a charitable deduction, (2) the charitable deduction as a percent of gross estates for those that claim the charitable deduction, (3) the ratio of legal, funeral, and other administrative costs to gross estate, and (4) the ratio of estate tax liability to taxable estate (the effective tax rate).

The four adjustments are then applied to our SCF predicted bequests. We first compute the SCF gross estate, which involves adding debts back to net worth (debts are a deduction in the published SOI tables). We then use the computed ratios for the given gross estate size class to subtract charitable contributions by probabilistically assigning a deduction using the fraction claiming, and the amount of the deduction using the ratio of charitable deductions to gross estates for those estates assigned a charitable deduction. Next, we subtract the funeral, legal, and other expenses as a fraction of gross estates. Then, we subtract the (known) SCF debts, solving for the taxable estate. At that point we apply the effective tax rate for the given gross estate size class, solving for estate tax liability. The net bequest to be distributed is then SCF net worth minus charitable deductions, legal and other costs, and estate tax liability. Finally, we divide the bequest by the number of children in order to generate an expected inheritance distribution.

For estates below the range covered by estate taxes—an increasing fraction because of rising estate tax thresholds—we make two simple adjustments when solving for bequests. The first is to assume \$10,000 (2016 dollars) in funeral/burial costs (based on data from the National Funeral Directors' Association). The second is to assume the greater of \$10,000 or 5 percent of gross estate for administrative costs, legal fees, realtor fees, and other deductions. The 5 percent number comes from inspecting the bottom of the estate tax range, where (for example) legal and administrative fees dominate charitable deductions as a fraction of gross estates.

Differential Mortality Adjustment

Demographic information about individuals is also used to compute the estimates of wealth that is (probabilistically) bequeathed, through a differential mortality adjustment. The mortality rates applied to SCF wealth holdings begin with SSA published cohort death rates by age and sex.⁴⁷ The published rates are measures of cohort death rates for given years (which we convert to cohort/age combinations) and we convert those to three year dates rates by first computing three year conditional survival rates then subtracting those from one.

However, there is a well-known problem with using average cohort/sex death rates by age generally referred to as “differential” mortality, meaning socio-economic status is negatively correlated with mortality. Failure to adjust for differential mortality would, in particular, generate too much in the way of estimated bequests at any point in time.⁴⁸

Congressional Budget Office (CBOLT) Mortality Differentials						
Control Variable	Male			Female		
	35-49	50-64	65-75	30-49	50-64	65-75
Income quintile 5	0.40	0.73	0.90	0.49	0.71	0.81
Income quintile 4	0.54	0.82	0.96	0.75	0.76	0.96
Income quintile 3	0.83	0.79	0.99	0.79	0.92	0.99
Income quintile 2	1.16	1.07	1.05	1.04	1.09	1.09
Income quintile 1	2.07	1.60	1.10	1.96	1.53	1.15
Less than high school	1.56	1.36	1.23	1.61	1.48	1.26
High school graduate	1.11	1.05	0.98	1.12	0.89	0.91
Some College	0.97	0.89	0.90	0.78	0.82	0.81
College graduate	0.55	0.64	0.62	0.58	0.64	0.68
Never Married	1.95	1.66	1.42	1.92	1.60	1.16
Married	0.72	0.85	0.90	0.75	0.81	0.83
Separated/Divorced	1.56	1.46	1.53	1.35	1.32	1.26
Widowed	1.53	1.93	1.26	1.53	1.44	1.16
White	0.90	0.95	0.98	0.93	0.93	0.96
Black	1.74	1.58	1.35	1.53	1.58	1.42
Other race	1.13	0.79	0.76	0.89	1.01	0.88
Hispanic	0.98	0.93	0.93	0.92	0.99	1.07

⁴⁷ The SSA data used is available at www.ssa.gov/OACT/HistEst/Death/2017/DeathProbabilities2017.html. The Stata code for computing three year death rates is available from the authors upon request.

⁴⁸ One other way to benchmark our estimates would be to compare the probabilistic bequests to actual bequests in estate tax data. However, substantial changes in estate taxes in the US in the time period we are considering has made it so very little of actual bequests will show up above the reporting threshold in recent years. In addition, differences in the valuation of various types of assets—especially closely-held businesses which are difficult to value—also makes a direct comparison with estate tax data problematic.

We adjust estimated bequests for mortality differentials using an approach developed by CBO for use in their long-term micro simulation model (CBOLT).⁴⁹ The table above shows CBO’s estimates of mortality odds ratios by age, sex, and four sets of demographic characteristics. A value of one in the table means that the specific demographic group has the same mortality as the average for the given age/sex group. Higher income, higher education, married, and white groups all have lower mortality than average. Importantly, all of the socio-economic variables found by CBO to be correlated with mortality differentials—income, education, marital status, and race—are also available in the SCF micro data as well.⁵⁰

A complication with simply implementing the mortality differentials is that CBO reports the odds-ratios independently across the four sets of socio-economic variables, so the four adjustments cannot be applied sequentially without adjusting for the fact that (for example) lower educated individuals also tend to be in lower income groups. In addition, since we use a different dataset from that which the CBO used for estimation, we cannot exactly match the odds ratios in the table above due to the variation in population weights within each category. We address this problem by first imposing that an individual’s mortality odds ratio is the linear sum of coefficients applied to dummies for each variable in the table above, and by making the following three assumptions: 1) The relative mortality odds ratios for income quintiles must match the table exactly, 2) within every other broad category—that is, education, marital status, and race—the *difference* between the mortality odds ratios must be maintained (i.e. for a male in the 34-49 age bracket, the difference in the mortality differential between a college graduate and a person who did not complete college, all else equal, should be 1.01, or 1.56-0.55), and 3) the population-weighted sums of the coefficients within the non-income categories should equal zero. We estimate these coefficients separately for each age bracket, sex, and year.⁵¹

The next step is to normalize these odds ratios such that the weighted average of *relative* mortality rates across the socio-economic groups have to sum to the overall cohort mortality rate (the SSA value) for any given age and sex group. We do this by scaling the odds ratios in each age, sex, and year group by a factor such that the average differential mortality equals the reported mortality rates reported by the SSA. In practice, this procedure produces univariate odds ratios quite similar to those in the CBO analysis, as shown in the table below.

⁴⁹ The CBO differential mortality model is described in Working Paper 2007-11, August 2007, available at <https://www.cbo.gov/publication/19096>.

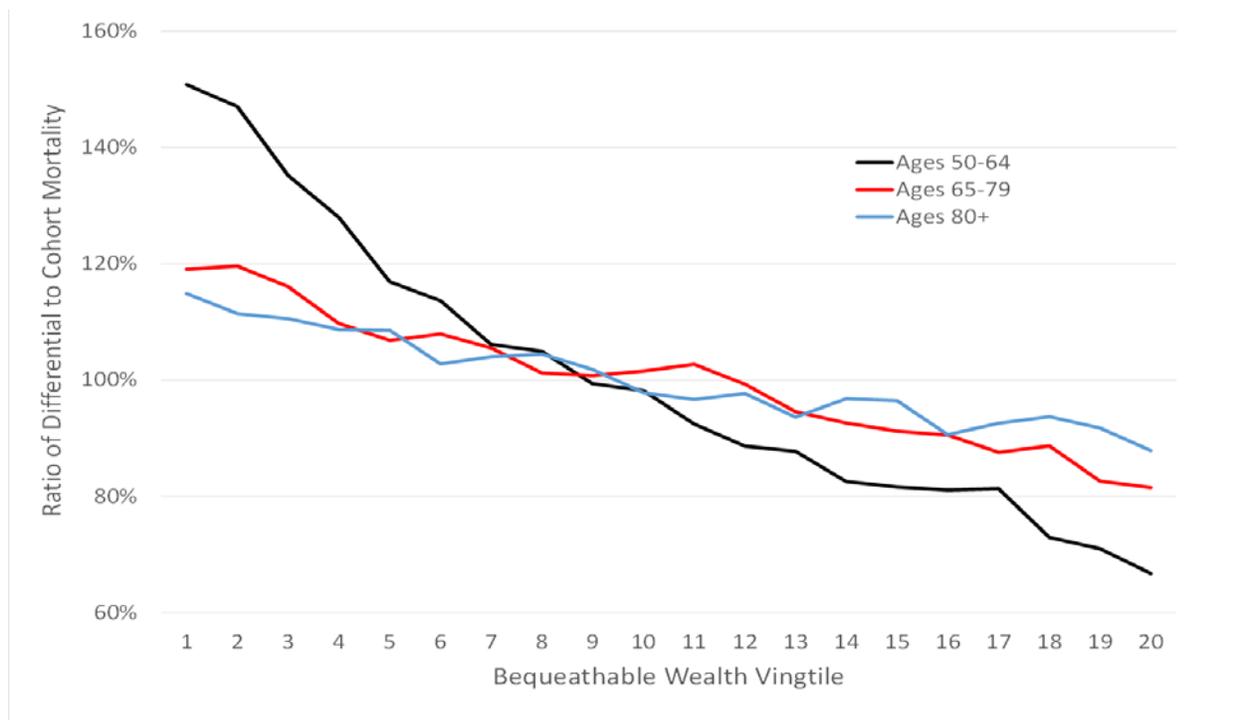
⁵⁰ The only notable difference between socio-economic measures is in the construction of the income variable. CBO’s estimates are based on SSA earnings and death records linked to various Survey of Income and Program Participation (SIPP) data files. The income quintile variables used by CBO are based (appropriately) on earnings over the lifecycle, not the SIPP current year earnings, in order to identify the “permanent” income that should in principle be what drives differences in mortality. To best match that, the SCF income concept used is the “usual” income variable collected in the survey, after the previous year income has been measured. Basically, respondents are asked if the income in the previous year is what they usually received, and if not, what that usual value is. For a further discussion of the usual income concept, see Box 4 in the latest (2017) *Federal Reserve Bulletin* article about the SCF, cited above. Finally, the concept of income in both the CBO estimates and the SCF are per-capita, meaning married couple incomes are divided by two before the quintile classification is assigned.

⁵¹ The Stata code for computing the differentials is available upon request. One pitfall of this linear estimation is that it does not constrain the resulting odds ratios for any particular person to be above zero. In practice, therefore, there are a few observations that fall in very low mortality types within the categories that have small negative values. We set these odds ratios equal to zero.

Simulated Mortality Ratios, Pooled over SCF years

Control Variable	Males			Females		
	35-49	50-64	65-75	30-49	50-64	65-75
Income quintile 5	0.43	0.71	0.80	0.53	0.68	0.72
Income quintile 4	0.56	0.79	0.87	0.76	0.72	0.89
Income quintile 3	0.84	0.78	0.98	0.78	0.91	0.98
Income quintile 2	1.17	1.07	1.10	1.02	1.09	1.15
Income quintile 1	1.98	1.64	1.24	1.90	1.58	1.26
Less than high school	1.58	1.46	1.42	1.71	1.63	1.46
High school graduate	1.17	1.14	1.06	1.24	1.04	0.98
Some College	1.05	0.96	0.90	0.87	0.94	0.79
College graduate	0.63	0.72	0.67	0.70	0.74	0.74
Never Married	1.81	1.57	1.33	1.71	1.48	1.20
Married	0.73	0.84	0.85	0.73	0.80	0.63
Separated/Divorced	1.62	1.41	1.18	1.39	1.27	0.90
Widowed	1.97	2.12	1.73	1.85	1.62	1.53
White	0.90	0.96	1.00	0.93	0.93	0.99
Black	1.70	1.54	1.15	1.47	1.54	1.24
Other race	1.05	0.74	0.70	0.85	0.90	0.64
Hispanic	0.90	0.88	0.82	0.92	0.96	0.94

As expected, implementing differential mortality introduces a negative correlation between mortality and wealth, because of the positive correlation between the socioeconomic variables and wealth. The chart below shows that the relationship between unadjusted and differential



mortality rates varies by wealth within three age groups. The higher ratios for lower wealth groups shows (for example) that individuals ages 50 to 64 at the bottom of the wealth distribution are almost twice as likely to die as those at the top of the distribution (relative differentials of 140 percent versus 70 percent). The gap shrinks with age (consistent with the underlying differential mortality inputs from CBO) but even among those 80 and older the mortality gap is around 50 percent.

The gross effect on expected bequests when shifting from unadjusted to differential mortality is about 20 percent, and consistent across the survey waves. Said differently, failure to adjust for differential mortality would lead us to *over predict* total wealth transfers made at death by about 20 percent. Having said that, much of that differential disappears when we look at net bequests *between* households, because more than half of death-related transfers in a given year go to spouses, and especially at younger ages where mortality differentials are most pronounced.

Measuring Inheritances and Substantial Gifts Received

The starting point for capturing inheritances and substantial gifts received is the inheritance module that comes near the end of the SCF survey.⁵² There is also (in some cases) supplemental information in the survey modules on real estate and owned businesses which both come in the early parts of the survey. In the real estate and business modules, respondents are asked how they obtained ownership of the asset, as part of the standard question battery, with “received as a gift or inheritance” as one of the options. In some cases those transfers are not captured again (as they should be) in the inheritance section. In addition, the question about “other” income in the SCF income module allows respondents to report an inheritance (cash or other financial assets only) received in the year preceding the survey year (to coincide with the timing of all other forms of income in the income module). Our comprehensive estimates of inheritances and substantial gifts received rely on information from all three parts of the survey.

SCF Primary Inheritance and Gifts Received Module

The inheritances module has retrospective questions on lifetime transfers received, with up to three occurrences for which details are collected, and a “mop-up” question to capture all other transfers received. Respondents are asked to report any inheritances or “substantial assets in a trust or other form” that they “ever received.” The data collected on the first three inheritances includes type of transfer, value of transfer, year received, and from whom. Note that there is no inquiry about what specific asset(s) were transferred, meaning distinctions like real estate versus stocks and bonds or cash.

The type of transfer variable is key for our allocation between inheritances and *inter vivos* gifts received. The type variable includes inheritances, trusts, and transfer/gift. In our analysis, the transfer/gift types are allocated to *inter vivos*, and the other types are inheritances received. The

⁵² The specific question wording used for all of the key variables described here is listed in section 5 below, along with the relevant possible answers (the “code frame”) when the answers are not dollar amounts or years. The question wording for the variables here has not evolved in the period we are using. The question wording and other key survey information for any SCF wave can be accessed at <https://www.federalreserve.gov/econres/scfindex.htm>.

other key variable in the inheritance module for our analysis is the year in which the transfer as received.

As noted, the SCF captures details (including type and year received) for the three largest inheritances or gifts ever received (beginning with the largest). The mop-up covers all remaining inheritances and gifts, and these are not negligible. For example, in the 2016 survey, the aggregate values for the first three lifetime inheritances received were \$4.2t, \$704b, and \$159b, respectively. The value of the mop-up inheritances or gifts ever received is \$272b, which is about 5 percent of the first three. The mop up is currently not included in our bequests received measures, because there is no obvious way to impute type or year received, but the fact that most of the measures we are working with only include inheritances and gifts received in the past three years suggests this is likely not a substantial omission.

SCF Real Estate and Business Modules

The biggest adjustments we make to the inheritances received estimates come from the real estate and owned business modules of the SCF, which appear near the beginning of the survey. For every type of real estate and business asset, there are questions about when and “how” the asset was obtained. In the case of real estate, the respondent is prompted by questions about what the asset was worth when it was obtained, and that cues a question for the interviewer to record the asset as having been purchased or received as a gift/inheritance. There is no distinction between gifts versus inheritances at the level of individual assets—they are lumped together as gift/inheritance. In the case of owned businesses, respondents are asked about sources of funds for investment in the business, with “inheritance” and “given” included in the code frame along with answers like “borrowed” and “used own funds.”

For both the real estate and business transfers, there should be a connection to the questions about inheritances and gifts received in the inheritance module (described above). In particular, if the respondent reports receiving a real property transfer in a given year, SCF protocol intends that there will be a corresponding inheritance or gift recorded for that year, though the specific asset may be a component of a larger reported inheritance that bundles multiple assets. Thus, the relationship is asymmetric, because a given inheritance may include both the real property being captured and other assets transferred at the same time. That is, reported inheritances should be at least as large as the real property received in a given inheritance year.

SCF cases undergo a rigorous review or “editing” process that captures and corrects many of these inconsistencies, but some do slip through, particularly in earlier waves when the editing software was less effective at capturing situations where (say) a respondent reported a house or business being received through inheritance but then failed to include that inheritance when they entered the inheritance module. Or, in some cases (and this is where timing within the survey matters) respondents may feel as though they are being asked a second time about the same event, even though the survey is trying to capture other details about that event, and (admittedly in some cases because they have been through a long and grueling interview already) are reluctant to answer the inheritance questions. This leads to situations where the sum of inheritances or gifts the respondent reports having received as real estate or businesses in a given year exceeds the total amount of reported inheritances and gifts in that year. In these cases, we

take on the incremental information from the real estate and business questions by marking the excess amounts as inheritances.

SCF Income Module

In addition to the adjustments for real asset transfers captured in the asset modules but missing from the inheritance module, there is a second type of potential inconsistency for “unaccounted” inheritances when the respondent reports income from an inheritance in the income section but fails to include that in the inheritance module.

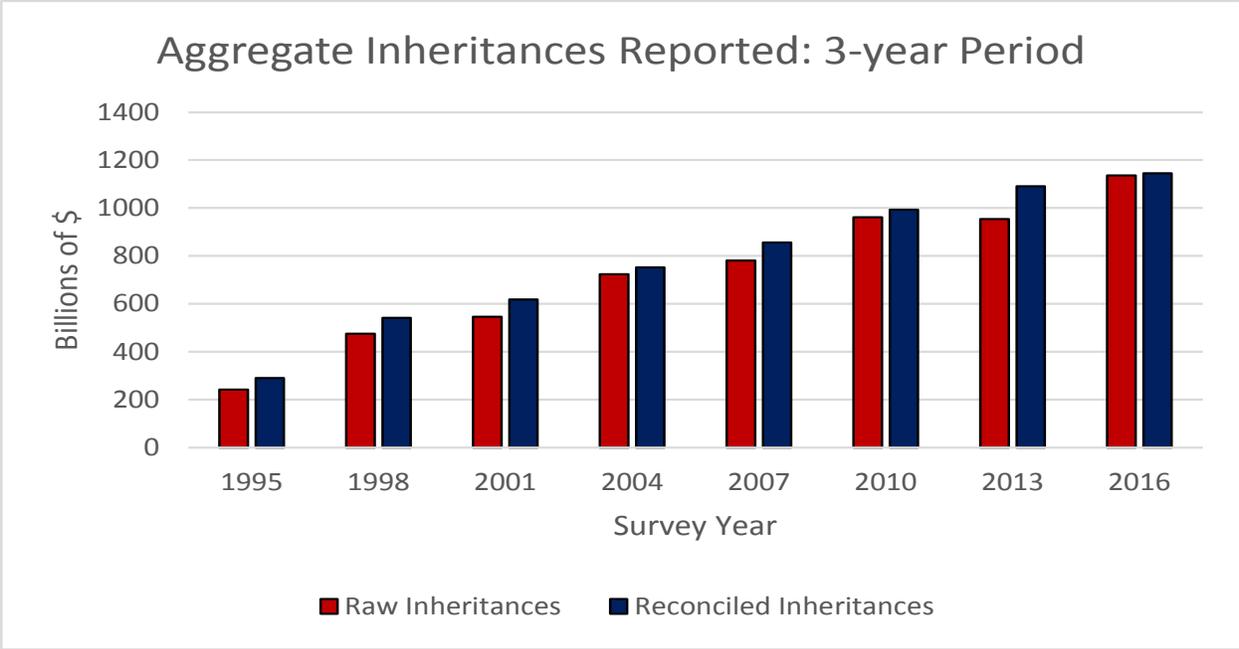
Inheritances of non-property (financial assets and cash) are included in the SCF income concept if received in the year prior to the survey, which is the same time frame for measuring the other income flows (and, we will see, most *inter vivos* transfers). The income section actually occurs before the inheritance section, so most often the inconsistencies arise because respondents do not include the financial asset inheritances and gifts as part of income, but then they report the prior year inheritance in the inheritance module. Thus, most of the case review/editing that occurs for this inconsistency goes the other way, meaning the reported inheritance is added to other components of income in SCF post-production.

However, there are cases that go in the other direction, where the income from an inheritance is reported in the income section but not in the inheritance section. As with real asset inconsistencies, we add those flows to the reported inheritances.

Reconciled Inheritances and Gifts Received

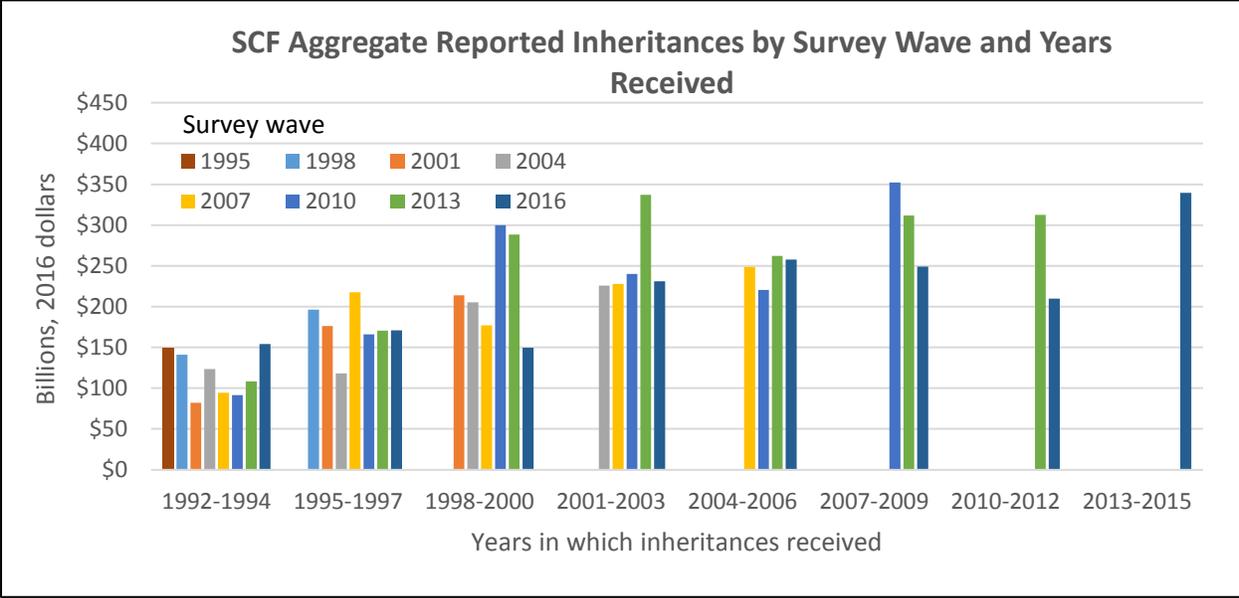
In practice, the primary inheritance questions do a good job of capturing all inheritances and gifts, and the adjustments we made for incremental information from the assets and income modules add relatively little to the total inheritance estimates. This observation in large part reflects the key SCF processing decision during case review: inconsistencies between the real estate/business and inheritance sections are flagged before the case is subject to review, and the inconsistencies are largely edited out before the final micro data is released to the public. The editing relies on interviewer notes as well as the underlying data itself.⁵³ The figure below shows the 3-year aggregates of both the “raw inheritance” aggregates—i.e. those estimated from the primary inheritance and gifts questions alone—and the “reconciled inheritances”, which are those that use the incremental information from the asset price questions. We use the “reconciled inheritances” for all of the estimates shown below.

⁵³ The inheritance module comes near the end of a long survey, and many respondents are rather exhausted. They may answer, for example, “I already told you about the inherited real estate.” The interviewer can make that note, keep the interview moving, and SCF staff then uses the information from the earlier module to fill in the missing inheritance.



“Look Back” Period for Measuring Inheritances and Substantial Gifts Received

Our top-line estimates for inheritances and substantial gifts received are based on a three-year look-back period, though for some questions we do rely on lifetime transfers ever received. Because the SCF inheritance module is retrospective over the respondent’s lifetime, we are able to compare the aggregate amount of inheritances received in a given observation period across survey waves to look for signs of reporting anomalies. In the chart below, we compare the amounts reported to have been received in all of the three-year periods covered by the survey waves that we use in this analysis:



There are two reasons why inheritances reported for a given three-year period may decline as we move further away from the observation period. First, some of the people who received an inheritance or substantial gift in a given observation period will have died before the next survey is conducted, so by definition their transfers received are not counted. Second, respondent recall about inheritances and gifts likely deteriorates with time. On the other hand, it is also possible that reported inheritances for a given time period will rise as we move further away from that time period, because of sampling variability and possibly because of improvements in how SCF cases are reviewed.⁵⁴

The chart confirms that in general aggregate reported inheritances and gifts do tend to decline as we move further away from the period for which inheritances are being measured. The left-most bar in each observation period is the survey wave that occurs right at the end of the observation period (consistent with a three-year look back period), the next bar to the right is for the survey wave at t+3 (looking back between t-4 and t-6), then t+6, etc. With a few exceptions, reported inheritances are at least as high in the waves closest to the three-year observation period as in other waves, and in a few cases (the period 2007-09 as captured in the 2010, 2013, and 2016 surveys, for example) there is a notable deterioration as we move further away from the adjacent three-year period. However, the observations are generally in the ballpark for all waves covering a given observation period, which suggests that recall and survivorship bias are probably not too large, which means that the part of our analysis based on lifetime recall is also robust.

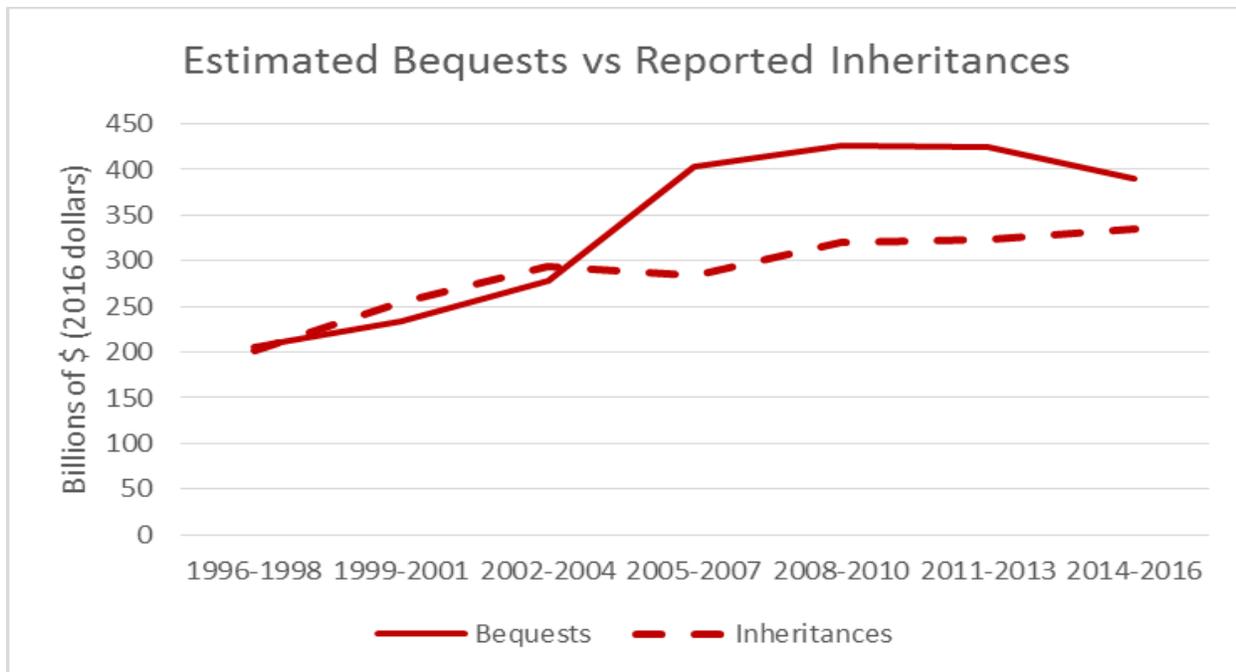
Comparing Estimated Bequests Made and Reported Inheritances Received

One goal of generating estimates of bequests made is to benchmark the reported values of inheritances received. The SCF is almost certainly the best micro data available for measuring inheritances received, because of the underlying sampling strategy (making sure high wealth households are included) and the substantial energy devoted to collecting inheritance information during the survey itself. However, there is still no way to know whether the inheritance amounts reported are reasonable; there are no administrative data against which to compare the estimates, except for the very high end where estates are subject to tax, and as noted above, even those comparisons are fraught with difficulty because of asset valuation considerations. In this section we compare reported inheritances against our estimates of bequests made using two tests. We look at the aggregate amounts given and received across three-year time periods, and the distribution of amounts bequeathed (adjusting for the number of likely recipients) and received within time periods. Both exercises are consistent in showing that the SCF does a very good job capturing inheritances received.

Summary statistics on bequests made and inheritances received for every three-year period between 1996 to 1998 and 2014 to 2016 are generated using the methods described above. Estimated bequests made are computed using the method described in section 2 (with a three-

⁵⁴ As described above, SCF case review protocol calls for reconciliation of reported asset transfers (real estate and businesses) and the inheritance and gift module. Many respondents report having received property but then do not report that as an inheritance or gift, and the case is edited to correct for that discrepancy, a process that also often relies on interviewer case notes. All of this was done by hand in early waves of the SCF, and has become increasingly automated and more efficient in recent waves.

year cumulative mortality rate) applied to the SCF survey wealth at the beginning of the period. Reported inheritances received over the three year period are captured by the SCF survey conducted at the end of the period using a three-year look-back window. Both sets of estimates are annualized. The time series results are shown in the figure below:



Keeping in mind that the sources of these two series are independent from one another and many assumptions go into the bequest estimates, the similarity in levels and trends suggest that the aggregate inheritance flows are well captured in the SCF.

Given the focus of this research on explaining the concentration of wealth, it is even more important that we capture the distribution of inheritances received, and not just the aggregate totals. The univariate comparison of the two distributions requires one additional assumption on the bequest side. For any given estate, we need to know how many potential inheritances are generated when the individual dies. The SCF has data on the number of living children for each respondent, so we use that variable to divide the estate into equal size potential inheritances. If the number of children is zero, we leave the estate as one large bequest.

Although the additional assumption adds yet another confounding factor to the bequest predictions, the counts and dollars transferred in each bequest/inheritance size class suggests that our approach is overall very effective, and (again) confirms that the SCF is doing a good job capturing inheritances received (as shown in Table 2 in the main text). The overall counts and aggregates across the period 1995-2016 show that the generally close relationship between the dollar value of bequests made and inheritances received reported above—with bequests slightly higher on average—also holds for the counts of bequests made and inheritances received (2.0 million bequests made versus 1.7 million inheritances received per year). The distributional statistics are also very reassuring. Both sides of the bequest/inheritance suggest a large portion

(36 and 40 percent) of inheritances occur at levels of \$1 million or above.⁵⁵ Although half of all inheritances and bequests are in the size range below \$25,000, both perspectives agree they account for only 5 to 6 percent of total dollars transferred at death.

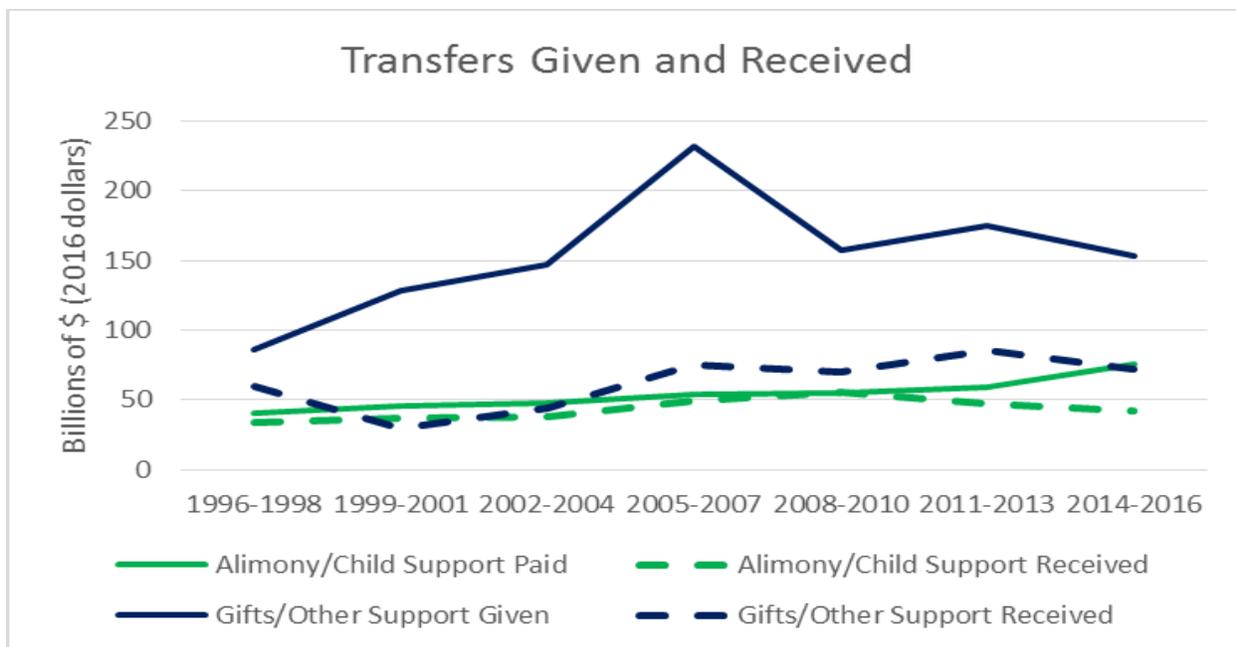
Reconciling Inter Vivos Transfers Made and Received

The other components of the two-by-two interfamily transfer matrix are *inter vivos* transfers made and received. In principle, “substantial” *inter vivos* gifts received are captured in the inheritance section, as described above, so we can add those to other types of *inter vivos* transfers received (mostly income support from others) and thus capture all transfers received. Using another SCF question on substantial gifts made and support paid to others, we (again, in principle) can see *inter vivos* transfers from the giver perspective as well. However, conceptual differences between gifts and support made and received lead to divergence in the aggregates as well as difficulties with separating the flows from other types of support given and received. In this section we show that the divergence between gifts made and received is quite large, and discuss what that means for tracking interfamily transfers more generally.

The SCF income module has two questions about income received (in the past year) that bear directly on transfers received. The first is about alimony and child support received. Alimony and child support received is a component of SCF *Bulletin* income, and runs about \$50 billion per year in recent waves. The second income module question involves the residual “other” income question. Two of the “other” income types (see the code frame in section 5 below) are for “other help/support” and “gifts, n.e.c.” Together, these amount to about \$20 billion in recent years, increasing from about \$6 billion in 1995. Anecdotally (based on interviewer comments and case review) this is probably a lot of parental support for adult children, but one cannot rule out that it also includes some misplaced alimony or child support.

The income module follow-up has two questions about transfers made to others, covering first alimony and child support paid, then other support paid and substantial gifts made. Separating alimony and child support paid from other gifts made allows us to compare alimony and child support flows head to head, which we do (green lines) in the chart below:

⁵⁵ We also ran the following thought experiment. What if all expected bequests were assumed to go to only one recipient? The number of expected bequests made falls by more than half, well below the number of inheritances received, and the distribution of expected bequests shifts wildly, with about 60 percent of the dollars showing up in the \$1 million or higher category, which is much higher than the inheritance received share. The experiment underscores that the expected bequest distribution is sensitive to how we assume estates are divided, and puts the differences between our baseline numbers (36 percent and 40 percent above \$1 million) in perspective.



In general, the data suggest that alimony and child support are well captured, with the perspectives of payers and receivers well aligned (a univariate distribution, not shown, tells the same basic story). Again, we cannot be absolutely sure that some of what is being captured in those questions does not reflect other types of *inter vivos* transfers, or that the question about regular support paid might include some alimony and child support, but the separation of those flows from the other types of *inter vivos* transfers is certainly plausible. Given that, the “other” income captured in the income module is then plausibly the income transfers we would expect (along with substantial gifts received from the inheritance section) to line up with *inter vivos* transfers made.

The instructions for the second income module follow-up question on *inter vivos* transfers made explicitly tell the interviewer to make sure the respondent includes “substantial gifts” that they made to others (excluding charities). The total reported transfer amounts are quite substantial, increasing from about \$60 billion in 1995 to nearly \$160 billion by 2016 (solid blue line in the chart above). This question provides the givers’ perspective on the gifts received that are captured in the inheritance section, but also includes other types of support paid, so it will be larger than gifts received. Indeed, the transfers made should be directly comparable to the sum of gifts received as measured in the inheritance section (which range from \$30 billion to \$50 billion over time) and “other” support income measured in the income module (which, as noted above, is now running about \$20 billion per year, increasing from about \$6 billion in 1995). The sum of those two flows is shown using the blue dashed line in the figure above, and is generally about half the amount of reported *inter vivos* transfers made.

The gap between *inter vivos* transfers made and received is substantial, though consistent with our priors because we expect that many such gifts and transfers reported being made in the SCF will not (and should not) be reported as being received in the SCF. One common example makes the point: a parent helping to support a child living outside the household and attending college will likely think of the costs of tuition, room, and board as “regular support” paid, while the child

on the receiving end will likely not report that as “income” received (though the survey would try to collect a regular parental cash stipend as part of “other” income). The distributional implications of these allocation decisions are shown in Tables 3 and 4 of the main text.

In on-going work we are exploring methods for allocating the unmeasured transfers received, assuming that the transfers made are accurately reported. The approach will ultimately involve delving deeper into the structure of the SCF, because we know (for example) the relationship of the person to whom the gift was made and (if a child) the age of that child. We also know (from the inheritance module) the respondent’s relationship with the person from whom the gift was received. Using these facts, combined with what we know about the univariate distributions of *inter vivos* transfers made, we can allocate the flows to transfers received.

SCF Question Wording and Code Frame Details

Inheritances and Gifts Received

The primary inheritance section, which comes at the end of the survey, asks:

X5801 Including any gifts or inheritances you may have already told me about, have you (or your {husband/wife/partner/spouse}) ever received an inheritance, or been given substantial assets in a trust or in some other form?

IF YES: Please do not include inheritances from a deceased spouse.

X5802 How many of these have you (or your {husband/wife/partner/spouse}) ever received?

#1 refers to the first gift/inheritance
#2 refers to the second gift/inheritance
#3 refers to the third gift/inheritance
#4 refers to all remaining gifts/inheritances

X5803(#1) Was that an inheritance, a trust, or something else?

X5808(#2)

X5813(#3)

1. *INHERITANCE; life insurance; other settlements
2. *TRUST
3. *TRANSFER/GIFT
6. *INHERITED TRUST
- 7. *OTHER

X5804(#1) What was its approximate value at the time it was received?

X5809(#2)

X5814(#3)

X5805(#1) In what year was it received?

X5810(#2)

X5815(#3)

X5818(#4) How much altogether were any others you have received?

There is also a possible entry for current year inheritances in the "other income" question in the primary income section:

X5723 (Other than withdrawals from account-type pensions or IRAs you told me about earlier in the interview, did/Did) you (or anyone else) have income from any other sources?

X5725 What other sources?

12. Inheritance/gifts

For housing and other real estate assets, the sequence about each property includes questions about the value of the property when it was acquired, and if the R indicates it was a gift or inheritance, that variable is checked. In the owned business section, the R is asked:

X3108(#1) How did you (or your family living here) first acquire this

X3208(#2) business; was it bought or invested in, started by you, inherited, given to you, or some other way?

1. *BOUGHT/INVEST
2. *STARTED
3. *INHERITED
4. *GIVEN
5. *JOINED/BECAME PARTNER/PROMOTION
10. Bought/invest and inherited
- 7. *OTHER

Regular Support Income Received

The primary income section includes the following questions:

X5717 Did you (or anyone else) have income from child support or alimony which you (or your family here) received?

X5718 In total, what was your (family's) annual income from child support or alimony which you (or your family here) received in {Survey Year -1}, before deductions for taxes and anything else?

X5723 (Other than withdrawals from account-type pensions or IRAs you told me about earlier in the interview, did/Did) you (or anyone else) have income from any other sources?

- X5725 What other sources?
13. Other help/support from relatives
28. Gift or support, n.e.c.

Alimony and Child Support Paid

After the income section, there are some additional questions capturing transfers made:

X5731 During {Survey Year -1}, did you (or anyone in your family living here) pay any alimony, separation payments, or child support?

X5732 Altogether, how much alimony and/or child support did you (and your family) pay in {Survey Year -1}?

Other Support Paid and Substantial Gifts Made

Continuing after the alimony and child support paid questions, there is another round of questions to capture other types of support and substantial gifts:

X5733 During {Survey Year -1}, did you (or anyone in your family living here) provide any (other) financial support for relatives or friends who do not live here?

Please do not include alimony or child support.
INCLUDE SUBSTANTIAL GIFTS.

X5734 How much support did you (and your family) pay?