

Lifecycle Patterns of Saving and Wealth Accumulation

Laura Feiveson
John Sabelhaus

July 2019

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 wealth change decomposition by age and agent type provides a new set of benchmarks for heterogeneous agent macro models, reconciling observed anomalies about lifecycle saving behavior and emphasizing the importance of generally unmeasured incomes (interfamily transfers and capital gains) in wealth accumulation dynamics.

Keywords: Household income, consumption, saving, wealth

JEL Codes: D14, H55, J32

Laura Feiveson is Principal Economist and John Sabelhaus is 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, Scott Baker, Jesse Bricker, Andrew Figura, Kevin Moore, and Jeff Thompson for detailed comments on an earlier draft, and to Elizabeth Holmquist for assistance with the Financial Accounts balance sheet and capital gains data. 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: (Corresponding author) Laura.J.Feiveson@frb.gov, jsabelhaus@gmail.com.

1. Introduction

Despite an enhanced focus on rising wealth inequality in the U.S. and other advanced economies, the role of saving behavior in wealth accumulation over the lifecycle is still an open question. This is because estimating the joint distribution of income, saving, and wealth 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, saving (or consumption), and wealth. 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. We overcome this hurdle by synthesizing two decades of micro data from the triennial Survey of Consumer Finances (SCF) with macro data from the Financial Accounts of the United States (FA) and the National Product and Income Accounts (NIPA). The SCF captures the top of the wealth distribution using a sampling and validation approach based on administrative data, and includes direct estimates of disaggregated balance sheet components, measures of interfamily transfers, labor incomes, and key demographic variables.¹ Using the synthesized micro data, we create pseudo-panels to study lifecycle patterns of wealth accumulation and consumption of U.S. consumers. By design—and in contrast to other lifecycle patterns estimated with microdata in the literature—the lifecycle patterns that we produce are consistent with FA and NIPA aggregates when summed over all consumers.

The first important data innovation required to build the pseudo-panel is to conceptually reconcile and then appropriately scale or impute the SCF wealth and income variables such that the aggregated micro data matches the FA and NIPA. We show that the reconciled SCF micro data generally line up very well with many of the FA and NIPA income and wealth aggregates, such that for most components we can simply use proportional scaling to reproduce the aggregate intertemporal budget constraint precisely. There are three wealth components—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 the cumulated capital gains on those assets. Thus, in our decomposition of wealth change using the

¹ For a description of the latest SCF results and a discussion of the administrative data sampling and validation, see Bricker et al. (2017).

SCF, saving summed across agent types matches the published aggregate, 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 over the lifecycle.

The constructed pseudo-panels make it possible to study the joint distribution of income, wealth, and consumption over the lifecycle and for various agent types. We start by disaggregating the sources of wealth growth by age and agent type into three components: conventionally defined (NIPA or 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, but net interfamily transfers are also key to positive wealth change at the end of the lifecycle. We then use the intertemporal budget constraint to solve for consumption, because income less consumption is equal to the change in wealth less capital gains and net interfamily transfers. Thus, although there is no direct measure of total consumption in the SCF, we are able to solve for consumption by age and agent type using the intertemporal budget constraint.²

Given estimated lifecycle flows, we then benchmark those flows using two different sets of ratios. The ratio of saving to disposable income is our first measure of the saving rate, because it corresponds to 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 rate is

² 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 the errors do vary with income and over the business cycle.

high early in the life cycle, and turns negative around typical retirement ages.³ The second way to characterize 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 many empirical studies, including ours, have found that wealth does not decline at older ages: capital gains and net transfers received by surviving agents at older ages are more than enough to offset negative conventionally defined saving.

The lifecycle patterns of wealth accumulation that emerge from the pseudo-panel disaggregation by agent type provide new insights about heterogeneity in U.S. saving and wealth accumulation. Saving, capital gains, and net interfamily transfers all play important roles in determining wealth change at various points in the lifecycle, but the relative magnitudes clearly differ by education and permanent income. Lower education and lower 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 modest wealth owned by lower education and lower permanent income agents is mostly in the form of housing, and much of the growth in that wealth component during the past two decades is because of house price appreciation. The highest permanent income group exhibits the highest saving rates at younger ages, well above the middle and lower income groups. However, negative saving at older ages holds for all agent types, and the rising ratio of capital gains on accumulated wealth to income by age is key to understanding the rapid wealth growth of the highest permanent income and education groups (relative to income) 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 is well described by Gabaix et al. (2016), Benhabib, Bisin, and Zhu (2015, 2017), Benhabib and Bisin (2018), and Benhabib, Bisin, and Luo (2019). However, there are open questions about how any given combination 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

³ The conceptual inconsistencies in cash flow saving estimates are mostly due to the treatment of retirement income. In particular, 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.

current versus future consumption in order to generate realistic wealth distributions. Some direct empirical analysis, such as Fagereng, 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.⁴ Some models such as Castañeda, Díaz-Giménez, and Ríos-Rull (2003), De Nardi, Fella, and Pardo (2016), De Nardi and Fella (2017) focus on non-standard stochastic labor income processes to solve the wealth concentration puzzle. Although we find strong evidence of heterogeneity in savings behavior, our results are consistent with the idea that saving behavior relative to conventionally defined income will never fully explain wealth concentration, because the fraction of wealth change explained by saving is a relatively small component of wealth change. Furthermore, since gains are such an important factor in wealth accumulation, it is imperative to understand portfolio decisions as well.

Our estimated lifecycle consumption profiles are also informative about differences in inequality based on various data sets and for various measures of economic well-being. 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 are 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 estimated differences in (for example) consumption across agent types are due to the population coverage, conceptual, and other measurement problems described by Attanasio and Pistaferri (2016).⁵ 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 high quality micro and macro data.⁶ Using the synthesized data, we estimate consumption for high earners that is well above estimates based on household spending data such as the Consumer Expenditure Survey.

A final contribution of the paper is to provide a framework for improving the understanding of key empirical joint distributions that are important for economic policy and forecasting. Disaggregated data on income, consumption, and wealth across agent types is often used to gauge differences in behavior at business cycle frequencies, generally by estimating

⁴ 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.

⁵ The problems with measuring income and consumption are also highlighted by Bosworth et al. (1991).

⁶ 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.

reduced-form concepts such as marginal propensities to consume out of income and wealth across different types of consumers, and how that affects macro outcomes in a model with heterogeneous agents. For example, Kaplan, Moll, and Violante (2018) focus on how illiquid asset holdings for a wide swath of the income distribution with considerable purchasing power affects consumption responses to income shocks, while Auclert and Rognlie (2018) focus on the role of differential shocks to (and thus differential uncertainty about) labor income. Although the empirical estimates suggests that this type of heterogeneity has relatively small direct effects, calibrated models suggest that the general equilibrium effects may be sizable. Such analysis is beyond the scope of this paper, but the pseudo-panel data approach here can in principle be used to directly test the implications of heterogeneity over the business cycle.⁷

The rest of the paper is organized as follows. In section 2 we introduce our intertemporal budget constraint accounting framework, focusing on the micro/macro data synthesis and interfamily transfers needed to disaggregate wealth change and solve for consumption 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 household 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 shows how we can reconcile basic lifecycle theory with the empirical data, and resolve long-standing misconceptions about lifecycle saving and wealth accumulation. In section 5, we show the lifecycle wealth change decomposition by agent type, linking groups across the cross-sections using education and relative rankings of permanent income within cohorts. Section 6 concludes.

⁷ The pseudo-panel approach is directly useful for studying more general ideas about heterogeneity and business cycles as well. 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 differed in the subsequent bust. 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. 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, because the rapid growth of mortgage debt and slowdown in spending (at least on new autos and housing) was most pronounced among higher (permanent) income groups.

2. The Intertemporal Budget Constraint in Micro and Macro Data

The textbook household intertemporal budget constraint is our starting point for measuring saving and wealth dynamics using the pseudo-panel decomposition. The simplest form of the budget constraint states that the change in wealth (ΔW) is equal to capital gains (G) plus saving (S), which is the same as disposable income (Y) minus consumption (C). Thus, we state the equivalence between the two measures of saving from either side of the budget constraint:

$$S = \Delta W - G = Y - C$$

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).⁸ The Financial Accounts (FA) balance sheet derivation of aggregate household saving begins with the left-hand side of the budget constraint, which is the change in wealth (ΔW).⁹ 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 the change in net worth over time. The FA also publishes the component of net worth change that excludes capital gains and is the same, in concept, to the NIPA concept of saving.

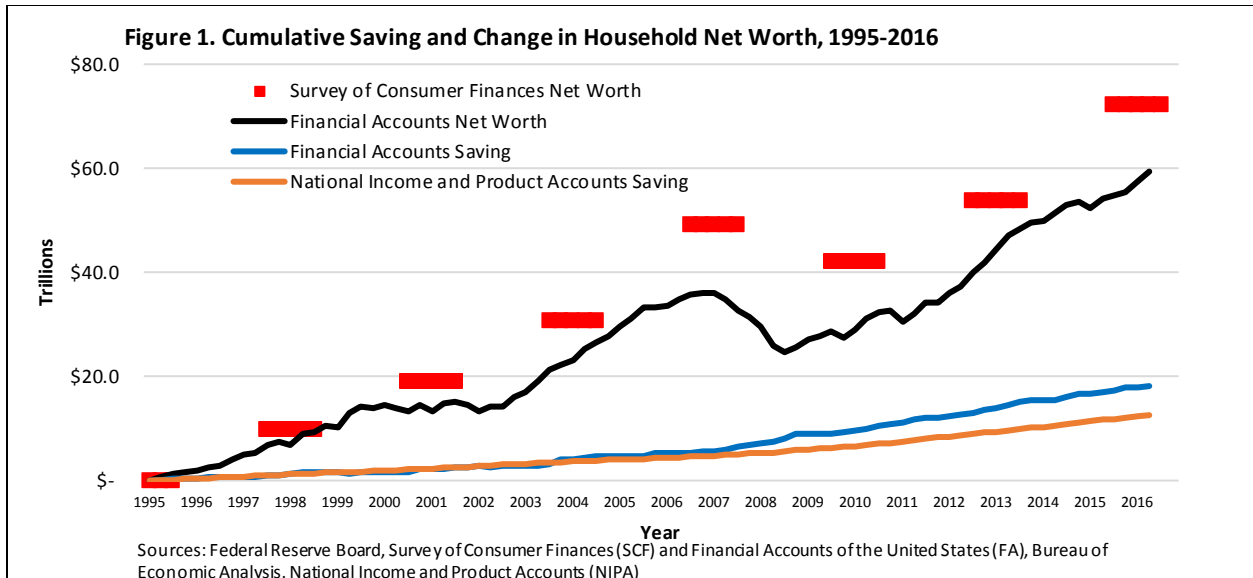
Saving and Wealth Accumulation in Aggregate Data

How much of aggregate wealth change can be accounted for by cumulated saving? Figure 1 shows two measures of cumulated wealth change and two measures of cumulated saving over the period 1995Q1 through 2016Q4.¹⁰ The two bottom lines show cumulated saving from the NIPA and the FA, and the two top lines show cumulated wealth change from the FA and the SCF. The key messages are that different data sources show varying magnitudes for the same economic concepts, but no matter how measured, cumulated wealth change over this period far exceeds cumulated saving.

⁸ See www.bea.gov/iTable/index_nipa.cfm.

⁹ The FA data is described in the Federal Reserve's Z1 release, see www.federalreserve.gov/releases/z1/current/.

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.



The cumulative change in FA household sector net worth is almost \$60 trillion for the 1995 through 2016 period, while the SCF-based change in wealth is higher, particularly after 2001 (more on this below). The two cumulated saving measures were in general agreement through the mid-2000s. In the second half of our sample period estimated aggregate saving in the FA has run slightly higher than in the NIPA. However, regardless of which wealth change measure is used and which cumulated saving measure is used, the following statement holds: saving accounts for less than a third of household wealth change during this period, meaning capital gains accounts for more than two-thirds of the total.¹¹ 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.¹² 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.

¹¹ The decomposition of wealth change in figure 1 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. Gale and Sabelhaus (1999) provide more details and a historical perspective on the theoretical and empirical relationship between FA and NIPA aggregate saving rates.

¹² Kuhn, Schularick, and Steins (2019) also focus on the role of capital gains in wealth accumulation, and show how the relative prices of different types of assets interact with the distribution of those assets to drive changes in wealth inequality. When house prices increased, the middle class benefitted the most. When financial asset prices increased, the gains were mostly at the top of the wealth distribution.

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.¹³ Household sector net worth in the SCF micro data grew much faster than the FA published aggregate over the 1995 through 2016 period. While FA aggregate household sector net worth (black line, figure 1) grew nearly \$60 trillion over the past two decades, the SCF (marked by the red squares spanning each SCF field period) growth was between \$70 and \$80 trillion.¹⁴ Given that the SCF is a survey with sampling and measurement variability, other research has suggested that the SCF is not properly capturing the value of key balance sheet components, and the solution has been to benchmark the SCF values (using proportional scaling) to the published FA aggregates.¹⁵

For financial assets and liabilities, the FA aggregates are derived from source data from financial institutions, and, like previous work, we deem that the FA is the appropriate benchmark. Thus, for these categories, we align the aggregate SCF level with the FA by scaling the individual holdings accordingly. Since the totals of these components already line up quite well between the SCF and the FA (see table 1), the scaling has only a small effect.¹⁶ We do not benchmark the SCF values to the FA for the other three categories: real estate, non-corporate business, and vehicles. These categories have difficult to observe market values, and thus the SCF values are more appropriate for our purposes, because they reflect what the respondents believe the assets are worth. The gap between SCF and FA equity in non-corporate businesses is due to both conceptual and measurement differences, but those are not easily disentangled.

¹³ See online appendix 1 for a detailed discussion of the steps taken here 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). Batty et al. (2019) describes the new Distributional Financial Accounts (DFA)—also based on reconciled SCF and FA balance sheets—now published quarterly by the Federal Reserve Board.

¹⁴ 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.

¹⁵ See, for example, Saez and Zucman (2016), Maki and Palumbo (2000), Sabelhaus and Pence (1999), and Cynamon and Fazzari (2016).

¹⁶ 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 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.

Balance Sheet Category	1995 Survey of Consumer Finances	Financial Accounts	
		1995 Q1	1996 Q1
Financial Assets	\$ 17.2	\$ 16.7	\$ 18.7
+ Real Estate	\$ 8.3	\$ 7.8	\$ 8.2
+ Noncorp Business	\$ 5.0	\$ 3.4	\$ 3.6
+ Vehicles	\$ 1.2	\$ 0.9	\$ 0.9
- Liabilities	\$ (4.3)	\$ (4.2)	\$ (4.6)
= Net Worth	\$ 27.1	\$ 24.6	\$ 26.8

Balance Sheet Category	2016 Survey of Consumer Finances	Financial Accounts	
		2016 Q1	2017 Q1
Financial Assets	\$ 59.3	\$ 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	\$ (13.4)	\$ (13.2)	\$ (13.7)
= Net Worth	\$ 104.5	\$ 80.4	\$ 87.1

Sources: 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.

The FA constructs the balance sheets of non-corporate businesses on an asset-by-asset basis, assigning market values to some assets such as real estate, for which price indexes exist.¹⁷ 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. 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 using published National Automobile Dealer Association (NADA) reports on a vehicle-by-vehicle basis. Finally, in the case of owned real estate, the FA is investigating a new methodology to value those assets, and implementing that change would eliminate much of the gap between FA and SCF housing values, raising the FA to be much closer to the SCF.¹⁸

The divergence between SCF and FA balance sheet aggregates is quite pronounced in these three categories, such that our choice not to benchmark them leads to a higher level of

¹⁷ FA Table B.104 shows the balance sheet decomposition for the non-corporate business sector.

¹⁸ 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.

aggregate wealth. 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. We do use the FA aggregate saving measure for those components, which means we are effectively assuming that the differences in within-category wealth changes over time between the SCF and FA are due to differences in perceived 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.¹⁹

Synthesizing Micro and Macro Incomes

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 our various saving rates across age and agent-type groups, and that requires income for the denominator. The steps needed to synthesize SCF micro and NIPA macro incomes are somewhat more involved than the steps for synthesizing balance sheets. First, the NIPA aggregates include many imputed components not available in the SCF. Second, the SCF only asks about incomes in the year prior to each triennial survey, so we (crudely) multiply the SCF values by three in order to compare to the NIPA across our seven three-year subperiods.

¹⁹ Bricker et al. (2016) directly assess how the decision to benchmark affects wealth concentration estimates.

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	\$ 12,088	\$ 11,278	107%
+ Business Income	\$ 2,175	\$ 2,370	92%
+ Social Security	\$ 675	\$ 1,047	64%
+ Interest and Dividends	\$ 1,245	\$ 3,216	39%
+ Other Government Cash Transfers	\$ 182	\$ 583	31%
+ Employer Retirement Contributions	\$ 273	\$ 833	33%
+ Retirement Interest and Dividends	\$ -	\$ 1,064	NA
- Personal Income Taxes	\$ 2,730	\$ 3,529	77%
- Employer Payroll Taxes	\$ 855	\$ 1,003	85%
= Adjusted Disposable Income	\$ 13,051	\$ 15,859	82%
Subperiod 2013 Q2 through 2016 Q1			
Disposable Income Components	SCF	NIPA	Percent Ratio
Wages and Salaries	\$ 24,087	\$ 22,906	105%
+ Business Income	\$ 5,196	\$ 4,839	107%
+ Social Security	\$ 2,247	\$ 2,554	88%
+ Interest and Dividends	\$ 2,023	\$ 5,220	39%
+ Other Government Cash Transfers	\$ 503	\$ 1,485	34%
+ Employer Retirement Contributions	\$ 579	\$ 1,496	39%
+ Retirement Interest and Dividends	\$ -	\$ 1,849	NA
- Personal Income Taxes	\$ 5,584	\$ 5,530	101%
- Employer Payroll Taxes	\$ 1,394	\$ 1,669	84%
= Adjusted Disposable Income	\$ 27,658	\$ 33,150	83%

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 disposable income concepts is available from the authors.

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.²⁰ SCF and NIPA incomes and taxes are in 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 2.²¹ The SCF directly captures the vast majority of corresponding NIPA incomes, but there is substantial

²⁰ See online appendix 1 for details about the steps taken to align NIPA and SCF income concepts.

²¹ 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.

variation across the income categories. In addition, for some income components the SCF has limited or even no direct information about the relevant income flows, but the wealth holdings underlying those income flows make it possible to distribute the incomes across cohort and agent type groups with great confidence.

The two largest sources of income are from wages and businesses, and for these there is a very close match between the SCF and NIPA aggregates.²² The third component of household income is interest and dividends on accounts outside the institutional retirement account system. In this category a combination of limited data and under-reporting leads to fairly large data gaps, but the SCF wealth data allow us to establish a solid distributional foundation for imputing and then scaling up to match the NIPA aggregates. One area where income information is very limited is for tax preferred financial assets outside the institutional retirement system (more on institutional retirement accounts below). These are assets in IRAs and annuities that generate income, but the associated incomes are generally not known by (or asked of) SCF respondents because they are not taxable. However, we use the balance sheet estimates of IRA and annuity wealth to assign income flows.

Measuring retirement income flows in institutional (defined benefit pensions and defined contribution account-type) plans using an internally consistent framework is important for our sources and uses disaggregation. Like most surveys, the SCF asks respondents about the incomes they actually receive as payouts from retirement plans, including both traditional pension plan benefits and withdrawals from account-type plans. However, the measure of income that is actually consistent with the intertemporal budget constraint is the new employer 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

²² Some observers, such as Bhandri et al (2019), have considered SCF business incomes and come to rather negative conclusions about the quality of the SCF data, focusing on the fact that reported business incomes in the SCF are much higher than the business incomes reported on tax returns. The SCF and NIPA totals line up well here because SCF business owners are likely reporting something closer to the economic income concept in the NIPA, and not the taxable business income concept in the tax data. This could be due in part to a willingness of business owners to report their actual business earnings on the survey, while they (or their accountants) report something different to the tax authorities.²² There is a sizable reporting correction in the published NIPA business income aggregates, which is why the SCF is simultaneously above the tax data and close to the NIPA. Other recent work by Smith et al (2019) shows that the very distinction between wage and business income is itself somewhat nebulous, and more of the labor income of business owners may (for tax reasons) be reported as business income in recent years, which could explain the trend in the ratios over our study period.

counterintuitive, the benefits paid and withdrawals from retirement accounts are not income, they represent the drawing down of accumulated wealth.

The SCF includes questions about both employee and employer contributions to retirement plans. Employee contributions are subtractions from wages and business income, so (assuming respondents report incomes before those deductions as the survey requests) those contributions are already captured as part of the underlying incomes (and remember, we match on wages). 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 SCF 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) institutional retirement accounts. However, the key to the validity of our benchmarking strategy is that there is no differential reporting of employer contributions to retirement plans across age or agent-type groups.

Lastly, in terms of retirement, the SCF has no data on the interest and dividends earned on institutional retirement accounts. These income flows are substantial, exceeding employer contributions and accounting for roughly one-fourth of all dividends and interest paid in the NIPA. 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, and again, especially for defined benefit pension plans. Respondents do have a good sense of the balances in account type plans (as described in the previous section) and we estimate the present value of future defined benefit payouts. Unlike taxable dividends and interest, however, the SCF does not even attempt to capture the information about capital income flows (table 2). Therefore, we allocate the missing interest and dividends using the reported retirement account balances and estimated present values for defined benefit pensions.

The under-reporting of government cash transfers in surveys is a well-known phenomenon. Social Security income reporting is much better, especially in more recent surveys where the information about benefits in the SCF labor force module are used to help inform (and impute as needed) the values in the survey income module. 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. Cash benefits are (by their nature) concentrated at the bottom of the income distribution at all ages, so proportional scaling of

reported benefits up is almost certainly unbiased. Likewise, Social Security has an obvious age component, and although there is some potential for bias if retirees who fail to report Social Security are at the bottom of the distribution, the effect is likely to be small.

SCF respondents are not asked about payroll or income taxes, but the information needed to fill out a tax form for respondents is collected in the survey. Thus, the two tax items are estimated using the NBER TAXSIM program, and the close correspondence between the SCF and NIPA tax aggregates is directly attributable to the close correspondence between the SCF and NIPA total taxable incomes and accuracy of the SCF distributions. One final adjustment is required due to the tendencies of businesses to underreport their income on tax forms. The business incomes of high-value businesses are adjusted in order to bring taxable SCF business incomes into line with tax data. With that adjustment, the ratios of estimated to NIPA income taxes are in line with the ratio of SCF to NIPA incomes.²³

Accounting for Interfamily Transfers

Interfamily transfers net to zero for the household sector as a whole, and thus those flows play no role in NIPA or FA intertemporal budget constraints. Accounting for such flows begins with a slightly expanded intertemporal budget constraint, which includes net interfamily transfers received (IFT) as a source of wealth change:

$$S = \Delta W - G - IFT = Y - C$$

Failure to account for interfamily transfers in a sources and uses disaggregation is potentially important, because in any given year such flows are roughly the same order of magnitude as the sum of all of household sector saving.²⁴ 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.²⁵ The important question when introducing transfers into the disaggregated intertemporal budget constraint identities is whether such transfers are measured

²³ Details on the use of TAXSIM with the SCF data are available on the SCF website.

²⁴ Feiveson and Sabelhaus (2018).

²⁵ 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.

well in the SCF data. That question is difficult to answer because, unlike income and wealth measures, there are no aggregate benchmarks against which to evaluate the survey data.

There are two principal forms of interfamily transfers. The first and largest form of interfamily transfer is bequests 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 flows from both giver and receiver is important in order to rearrange the disaggregated (birth cohort and agent-type) intertemporal budget constraints and thus disaggregate saving and consumption. The SCF survey instrument directly captures most of the flows required for the intertemporal budget constraint disaggregation (the survey sample frame does not include estates) and by estimating the missing flows and comparing transfers from both the giver and receiver perspective, we have a data check on the internal consistency of the estimated transfers.²⁶

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. In addition, we estimate bequests made using a model of differential mortality and adjustments for inheritance taxes, funeral expenses, and other death-related costs.²⁷ Also, the bequest made by a deceased individual does not have a one-to-one correspondence with reported 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.

Table 3a 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 transfer 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 \$287 billion per year from the inheritances side. Remaining divergence is due to unreported inheritances or misspecification in the bequest imputation.

²⁶ Online Appendix 2 provides details about the interfamily transfer measures described in this section.

²⁷ Consistent with the notation introduced in the next section, the SCF does not ask respondents about inheritances received by surviving spouses.

Table 3. Interfamily Transfers, 1996 to 2016				
a. Bequests and Inheritances				
	Bequests Made		Percent of Total 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
b: Intervivos Transfers Part 1: Alimony and Child Support Paid and Received				
	Support Paid		Percent of Total 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
c: Intervivos Transfers Part 2: Voluntary Gifts and Support Given and Received				
	Support Paid		Percent of Total 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 distributions of transfers at death by size from the giver (bequest) and receiver (inheritance) perspectives also 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 transfer 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.

The SCF survey instrument captures both sides of *inter vivos* transfers directly. Table 3b 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.

The final category of interfamily transfers is *inter vivos* gifts and support other than alimony and child support. Table 3c 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 3a 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 *inter vivos* transfers made information is collected after the income module, and 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.” The amount and number of gifts and support paid are higher than amounts received, because the transfers made is a much broader concept. There is evidence that large *inter vivos* transfers of greatest interest are captured 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 other 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 likely do not report those as transfers received. In the empirical work we calibrate the voluntary *inter vivos* transfers to match the amounts received.

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 and 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 the spouse/partner of the head of household.²⁸

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 i (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 across heads and their spouses/partners. 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 intertemporal budget identity—namely, by decomposing the change in wealth between savings, gains, and net interfamily transfers to that group. (In the aggregate, the transfers net to zero.) To back out savings from this identity, we need to use the SCF to get estimates of the size and

²⁸ 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.

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:

- 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.²⁹ 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 *inter vivos* gifts made and received) at the beginning of the three year period. This implies that the capital gains that accrues on that transferred wealth during the three year period will be credited to the group receiving the transfer at the beginning of the current three-year period.

Bequests Made

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,

²⁹ 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 with several roommates will be in the PEU.

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

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,

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

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 $W_{jct} = \sum_{i \in j,c} w_{it}$:

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

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 j and cohort c at time t through the direct spousal link is:

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

Where w_{it}^s is the wealth of the spouse of individual i , and j_s and c_s are the observed agent type and cohort of the spouse. The total pool of non-spousal bequests is given by,

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

These remaining bequests are distributed across all other surviving individuals.³⁰

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 $b^{+ns}(j,c,t,x_{it}) * (1 - d(j,c,t,x_{it}))$.

³⁰ 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.

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 $V_{jct(t,t+3)}^+$ as the *inter vivos* transfers received by the j, c group and $V_{jct(t,t+3)}^-$ 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. We assume that real estate and marketable financial assets all earn proportional gains, with a gain rate of $g_{(t,t+3)}^z$, over each three year period. For these assets, the capital gains earned by the survivors in group j, c between time t and $t+3$ is:

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

Capital gains for non-corporate businesses are assigned using respondent reported cost basis relative to current market value of the business. In effect, for any given group j, c , we observe their unrealized business gains (the gap between market value and cost basis) at both times t and $t+3$, thus we have a direct estimate of gains accrual.³¹

³¹ The approach is motivated by and consistent with Avery, Grodzicki, and Moore (2015), who compare and contrast traditional estate taxation versus taxing unrealized gains at death.

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. Our first approach is to compute and report per-capita values for the various measures across cohort groups, agent types, and sub-periods.

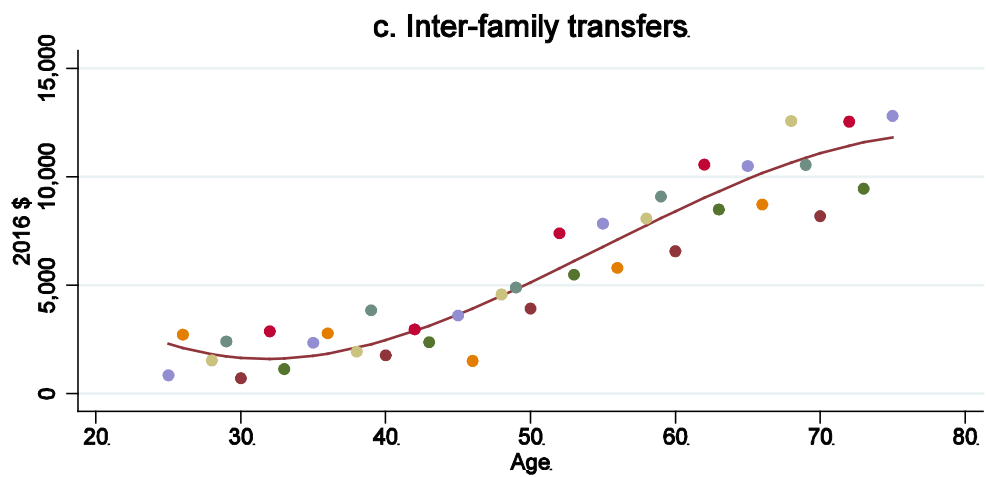
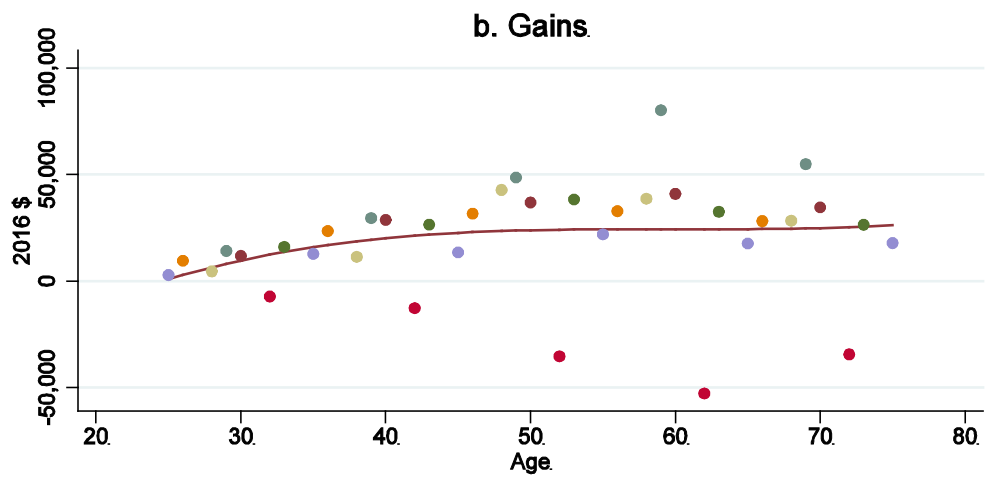
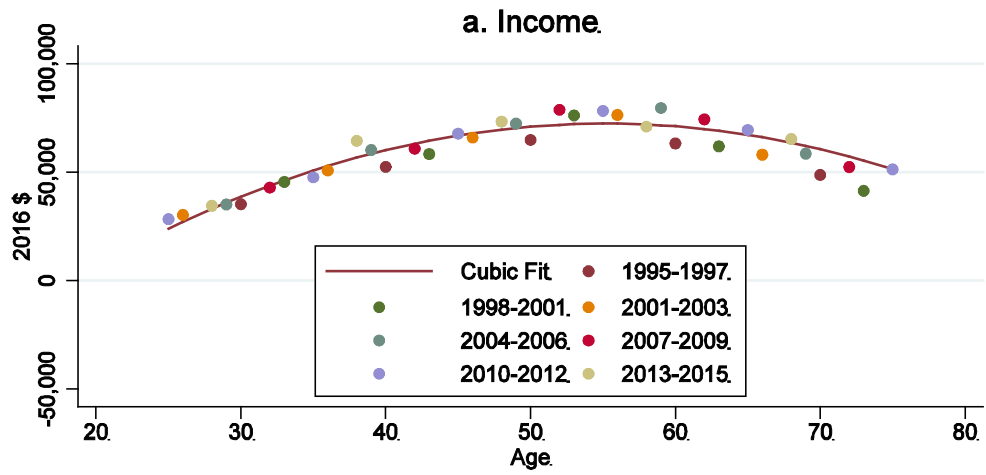
4. Lifecycle Saving and Wealth Accumulation

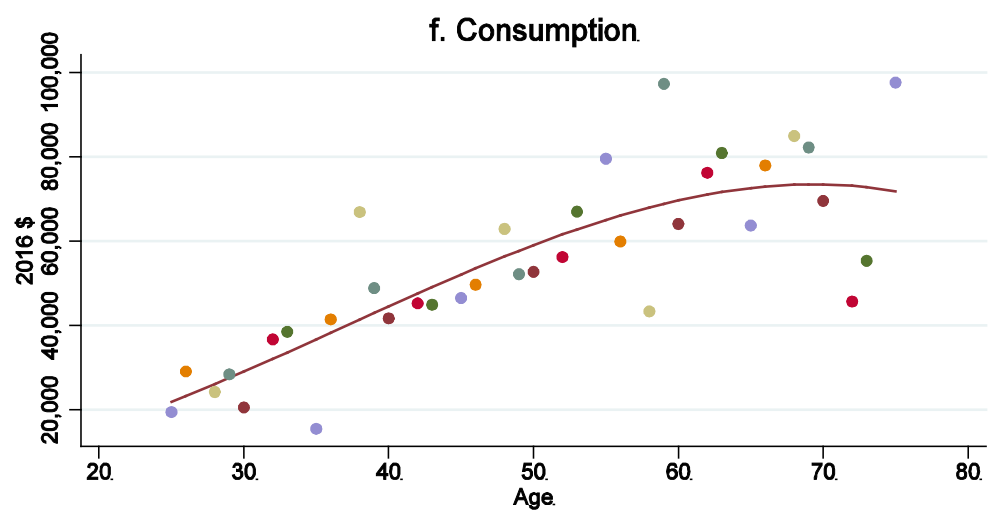
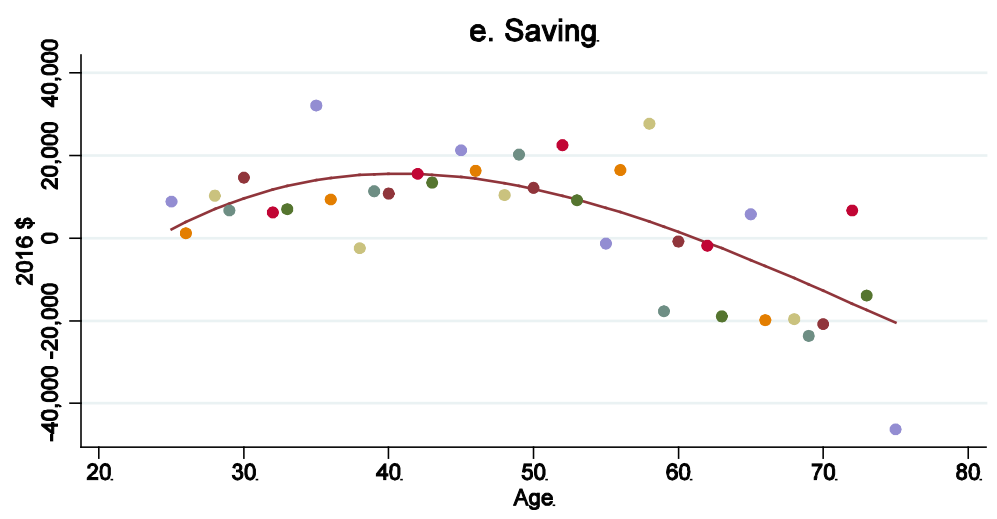
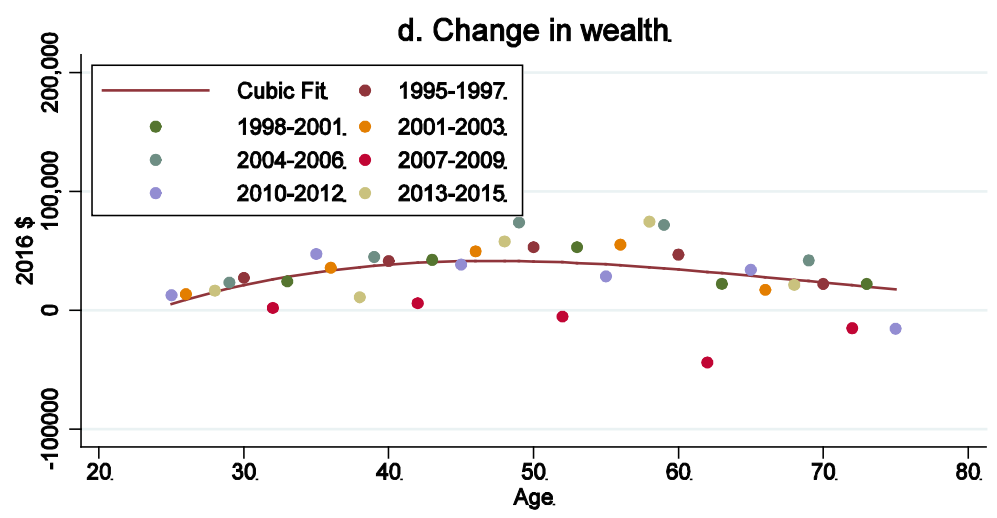
The pseudo-panel decomposition presented in the previous section provides the empirical framework one needs to construct estimates for lifecycle patterns of saving and wealth accumulation, but there is a statistical hurdle. The SCF is a relatively small sample, and although the sampling strategy carefully represents 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 are subject to 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 reconcile in other micro data sets and using alternative empirical approaches.

The first set of results (figure 2) are scatterplots of the sources and uses of funds for birth cohorts born between 1910-19 and 1980-89. The first three sub-charts show lifecycle sources of funds: disposable income, capital gains, and net interfamily transfers received. The second set of scatterplots shows lifecycle uses of funds: change in wealth, conventionally defined saving, and consumption. All values are per-capita for the head and spouse—meaning the household total is divided by two for households with a head and spouse, and unadjusted for single person households. The per-capita values are all in 2016 dollars.

Each birth cohort is represented as a series of point estimates color coded by sub-period, but not all cohorts show up in every one of the three-year SCF sub-periods. The oldest cohorts are not included at the end of the sample period because they are mostly deceased by 2013-2015. Similarly, the youngest cohorts are not included at the beginning of the sample period because they were not in their own households in the 1995-1997 subperiod. In total, there are 36 data points in panel of figure 2, 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. Thus, it is possible to track a cohort along the age axis because the median age for that cohort increases exactly three years between surveys.

Figure 2. Sources and Uses of Funds over the Lifecycle
 Per Capita for Household Head and Spouse, 1995-2016.





The line on each graph in figure 2 shows the smoothed lifecycle profile of that component when we pool over all cohorts and across the seven time periods. That the dots cluster around the smoothed line for many components at most lifecycle points—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 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 ensures 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.

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. In particular, per-capita change in wealth (panel d) is clearly increasing most rapidly through middle age, but, unlike the textbook model, wealth is still increasing (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 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 textbook lifecycle behavior.

The decomposition of wealth change into component flows using the intertemporal budget constraint in figure 2 helps make it clear why the data and the theory seem at odds. The lifecycle pattern of per-capita disposable income (panel a) has the usual hump-shape shown in textbooks, despite the potential bias from looking at multiple cohorts on the same chart.³² The other key ingredient in the textbook charts is negative conventionally defined saving at older ages, which shows up clearly in panel e. Average per-capita saving turns negative around 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

³² All of the per-capita values in figure 2 are in real terms, but younger cohorts have higher real lifetime earnings, and thus the true income trajectories for a given cohort are steeper and decline less at older ages.

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 defined saving such that per-capita wealth continues to grow.

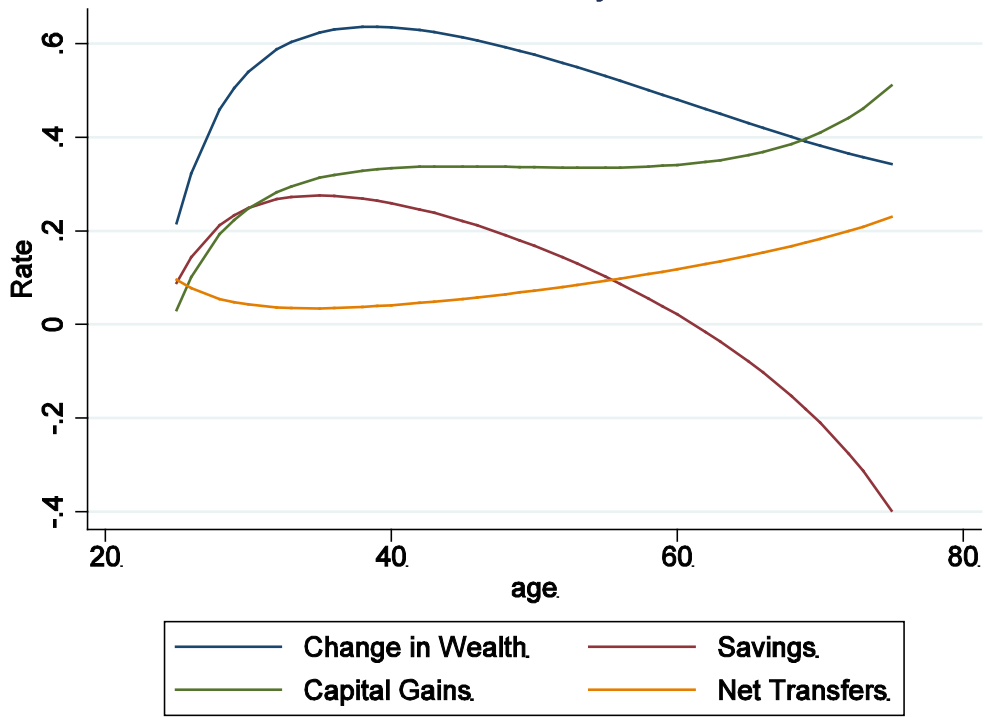
Our finding that retirees have negative saving on average is, at face value, in contrast to most of the literature that concludes retirees do not dissave. The finding in the literature that most retirees do not dissave is mostly driven by a definition of “income” that is inconsistent with the NIPA definition we use. While a retiree may think of a pension benefit or 401(k) withdrawal as income, it is actually a drawdown of accumulated wealth from a NIPA perspective. Our estimate of saving is consistent with the NIPA concept, which is why we do find negative saving rates in retirement. The fact that retirees do not spend their entire pension benefit just means they are dissaving more slowly, not that they are saving in a NIPA sense.³³ Similarly, at earlier stages of the lifecycle, the employer contributions to along with dividends and interest earned on pensions and other retirement accounts are a form of (generally unmeasured) 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 Fagereng, Holm, Moll, and Natvik (2018). Both of those papers struggle with the same problem raised here. Should the increases in wealth over the lifecycle associated with capital gains and interfamily transfers be included in measured saving, as intended additions to wealth? Our estimates of consumption, solved for using disposable income minus saving and shown in panel f of figure 2, indicate a steady increase in spending over the lifecycle with a leveling-off after retirement. Individuals are clearly increasing consumption as their income and wealth accumulate over the lifecycle, but not fast enough to spend down the increases in wealth.

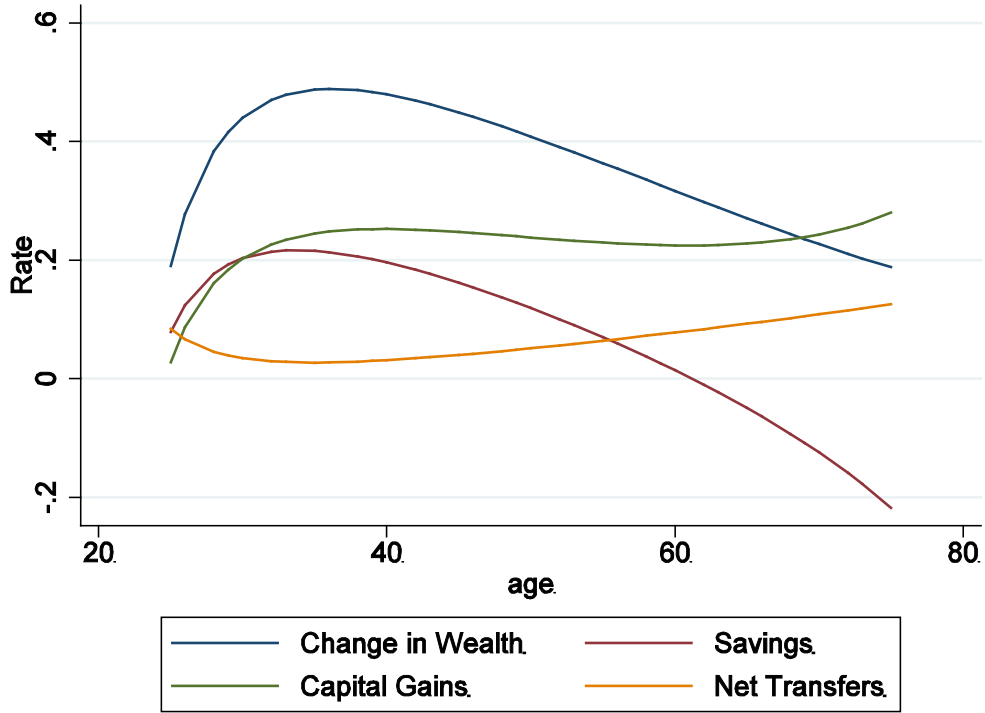
³³ Unlike private retirement wealth, Social Security does not have this accounting problem, because contributions are a tax, and benefits are a transfer, and both are accounted for in the NIPA disposable income concept. In principle, one could construct a measure of Social Security accrual and draw down that is comparable to the treatment of private retirement wealth.

Figure 3. Flows as rates.

a. Flows divided by income.



b. Flows divided by gross income (income plus gains plus transfers).



Some of that lack of spending is not difficult to understand, however. 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. A similar argument applies to a business owner whose labor income is directly tied to owning a business asset: the owner might not consume more when that asset goes up in value, if higher consumption means selling the asset and having to forego the labor income.

The connection between lifecycle behavior and our disaggregated sources and uses of funds is represented two different ways in the flows versions of the pseudo-panel charts (figure 3). In both panels, we consider the change in wealth and its three source components: saving, capital gains, and net interfamily transfers received. In the top panel of figure 3 we normalize flows by disposable income, and in the bottom panel, we normalize using the sum of all three sources of funds, which is disposable income plus capital gains and net interfamily transfers received.³⁴ The immediate effect of moving from the top to the bottom panel is on the range of the smoothed lines on the vertical axis, because the same numerators are divided by larger denominators. That effect becomes more important at older ages because gains and net interfamily transfers received account for a growing share of sources of funds.

The red line in panel a of figure 2 shows our pseudo-panel estimate of the conventionally defined saving rate over the lifecycle. By virtue of having synthesized the micro and macro data, this saving rate is the lifecycle analog to the NIPA saving rate over the 1995 through 2015 period.³⁵ The pattern is very different (to our knowledge) from other estimated lifecycle saving profiles, however. Most notably, saving is relatively high early in the lifecycle, steps down gradually as agents near retirement, and is decidedly negative at older ages. Again, much of the apparent discrepancy between this theoretically consistent saving rate and what is usually reported has to do with what is actually being measured, and how that relates to the intertemporal budget constraint (see section 2). Investments in housing, owned businesses, defined benefit pensions, and retirement accounts are key components of saving for individuals early in the lifecycle that are often uncaptured or only partially captured in household data. At the other end of the lifecycle, pension benefits and other retirement account withdrawals, sales of owned

³⁴ The lines in figure 3 are ratios of the fitted per-capita lines in the various panels of figure 2.

³⁵ The disposable income concept we are using here excludes some in-kind imputations, such as employer and government provided health care, so this saving rate is slightly higher than the published NIPA saving rate.

business, or sales of housing (with the rest of the balance sheet held constant) constitute dissaving, that is similarly unmeasured or counted as a source of income in other data sets.

The differences between the estimated flows in panels a and b are instructive. The numerators are identical, so (for example) conventionally defined saving as a flow still turns negative around age 60 when divided by income or the sum of income, gains, and net interfamily transfers received. However, panel b shows more rapidly declining rates of wealth accumulation and relatively flatter capital gains relative to the expanded denominator, as capital gains in particular become more important. These differences reinforce the idea that spending out of capital gains is likely low relative to other sources of funds. The implied reluctance to spend out of capital gains is consistent with direct evidence from Di Maggio et al (2018).

Viewing interfamily transfers in relative terms is also key to understanding increases in wealth and consumption at end of the lifecycle. Figure 3 shows that net transfers received account for an increasing share of and even dominates total wealth change at the oldest ages, regardless of which denominator is used. This observation helps explain how these synthesized micro/macro lifecycle patterns can be reconciled with the existing empirical literature on saving and consumption at the end of the lifecycle. The difference between household and per capita (head and spouse) as the unit of observation is crucial, because much of net interfamily transfers at death are between a head and spouse. One can control for such transfers using a household panel frame, but the other side of the intertemporal budget constraint shows us why in that approach household consumption declines at the death of a spouse.

There are two potential empirical issues likely affecting the estimated lifecycle patterns, both of which suggest our estimated consumption of the young is too low (meaning their estimated saving is too high) while consumption of the middle aged and old is too high (their saving is too low). First, there is evidence we are clearly missing *inter vivos* transfers made by older individuals (see section 2), and any missing transfers show up as consumption of the middle aged and old, instead of consumption of the young. Second, our assumption of proportional capital gains on non-business assets is biased if (say) young households invest more time in home improvement, and by virtue of that should receive disproportionately higher gains. These observations may help explain why, for example, we find relatively flatter consumption profiles in retirement than are suggested by household spending data.³⁶

³⁶ See, for example, Gourinchas and Parker (2002).

5. Lifecycle Patterns by Agent Type

The cohort-level analysis in the previous section offers some key new insights about lifecycle saving and wealth accumulation patterns, and in this section we use the pseudo-panel approach to disaggregate lifecycle flows by agent types within cohorts. Disaggregation is particularly important for answering the sorts of questions about wealth concentration posed in the introduction. Furthermore, calibrating heterogeneous agent models requires an empirical basis for introducing differences in behavior by agent type over the lifecycle. The disaggregated pseudo-panels also generate new findings about measured inequality, because unlike estimates based on other available micro data, our estimated income, consumption, and wealth change are internally consistent and sum to the macro aggregates.

An admissible agent type is based on any characteristic that divides the population in a unique way from one cross-section to the next. More specifically, agent types are constructed such that a specific member of a given agent type will not change to another agent type from one cross-section survey to the next.³⁷ Birth cohort (as in the previous section) is the most obvious pseudo-panel classifier, because every birth cohort is older by exactly three years between SCF survey waves. The two within-cohort agent types we consider in this section are education and permanent income, where permanent income is based on the SCF respondent-reported transitory versus permanent income decomposition.³⁸

There are tradeoffs when grouping by agent type, and we use the SCF oversampling strategy to our advantage to create sub groups that most effectively separate the top of the distribution from other groups in order to focus attention on wealth concentration. For education, we split the population into three groups: high school or less, some college, and college degree or higher (table 4). For permanent income, we split the population into the bottom 50 percent, the next 40 percent, and the top 10 percent (table 5). Due to the oversampling of higher-wealth individuals in the SCF, the top 10 percent group has nearly as many unweighted observations as the other two percentile groups.

³⁷ It is worth a reminder here that a key reason we estimate lifecycle patterns per capita for head and spouse/partner and not per household is that a given household is much more likely to “change” agent type from one survey to the next. Focusing on the head and spouse separately makes our agent type identification much more robust.

³⁸ 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 from backing out transitory shocks using panel income data. One important key to using the permanent income agent type is that agents are sorted by permanent within their own birth cohorts.

	High School or Less					
	<35	35-44	45-54	55-64	65-74	75+
Income	23,771	37,501	41,490	41,010	35,152	37,304
Capital Gains	3,553	7,166	9,161	9,645	9,715	10,736
Net Interfamily Transfers Received	769	1,001	1,953	3,929	6,143	8,197
Net Change in Wealth	6,976	10,693	13,664	9,011	3,125	10,456
Saving	2,654	2,526	2,550	(4,564)	(12,733)	(8,477)
Consumption	21,117	34,975	38,940	45,574	47,886	45,619
	Some College					
	<35	35-44	45-54	55-64	65-74	75+
Income	27,788	50,655	59,390	59,021	55,388	63,254
Capital Gains	4,998	12,231	19,477	17,160	17,677	31,941
Net Interfamily Transfers Received	1,617	1,852	3,167	6,868	9,656	15,822
Net Change in Wealth	11,264	19,928	27,099	27,550	9,476	13,527
Saving	4,649	5,845	4,454	3,522	(17,856)	(34,235)
Consumption	23,138	44,809	54,936	55,498	73,245	95,975
	College Degree or Higher					
	<35	35-44	45-54	55-64	65-74	75+
Income	39,695	83,021	113,694	128,791	109,791	96,232
Capital Gains	12,414	32,818	48,629	60,031	44,713	45,747
Net Interfamily Transfers Received	3,115	4,924	8,760	14,900	20,377	24,486
Net Change in Wealth	33,223	65,651	96,125	91,198	50,516	34,496
Saving	17,694	27,909	38,737	16,268	(14,574)	(35,737)
Consumption	22,001	55,112	74,957	112,523	124,365	130,076

Lifecycle Flows by Education

Table 4 shows lifecycle source and uses of funds disaggregated by education within six age groups. As in figure 2, all values are annualized per-capita in 2016 dollars, but in this case averaged over the entire sample period 1995 through 2016.³⁹ Education is highly correlated with lifetime income and wealth, and the connection is strong enough to draw a number of conclusions about how saving and wealth accumulation behavior varies over the lifecycle for different types of agents.

The starting point is to note that incomes are fairly close for the two lower education groups at youngest ages, while incomes of the college-educated are about 50 percent higher. By middle age the income gaps open up substantially, with the incomes of the some college group almost 50 percent higher than high school, and incomes of the college and greater group roughly

³⁹ The specific connection to figure 2 is that the cohort groups (or subset of cohort groups) that fall in a given age range are mapped into the appropriate column of table 4. Thus, as in figure 2, a given birth cohort may appear in two or even three columns, because they are tracked for as many as 21 years.

twice that of the some college group. In retirement, incomes in the bottom education group do not fall as sharply as in the two higher education groups, which reflects the fact that Social Security is the key source of retirement income for low lifetime incomes, while (again) receipt of pension benefits or withdrawal from retirement accounts are a drawdown of wealth. The differences in age-income profiles are otherwise consistent with observations from other income-oriented data sets. Most notably, both the level and steepness of the profiles rises with education.

Despite very different lifecycle income profiles, other lifecycle patterns share some important features across the education groups. Conventionally defined saving starts out as positive for all education groups, and turns negative at older ages. Capital gains also grow in importance over the lifecycle for all groups, though there are of course differences in the types of assets held by the three groups, with housing dominating the high school and some college groups, and financial assets and businesses more important at the top. Finally, net interfamily transfers received also grow in relative importance over the lifecycle for all education groups, though the estimated levels are of course very different.

Although conventionally defined saving starts out as positive and eventually turns negative for all education groups, the lifecycle profiles diverge by education over most of the lifecycle in predictable ways. Saving remains low and relatively flat for the bottom education groups, but continues to rise through middle age for the highest education group. Again, the intertemporal budget constraint approach underlying these estimates is useful for understanding the patterns. The methodology observes that cross section financial wealth holdings and claims to retirement wealth are highly concentrated in the top education group, and thus infers that (after subtracting capital gains) the highest education group saved substantially more.

The differences in lifecycle consumption profiles are consistent with the other lifecycle sources and uses, though the potential for bias from unobserved *inter vivos* transfers described above is likely disproportionately affecting the highest education group. For example, if parents help their children with the down payment on a house and we do not capture that flow in our net interfamily transfers, the increase in wealth is interpreted by the intertemporal budget constraint as the recipients' own saving out of current income, and given their observed income, that also implies lower consumption. Similarly, if the young higher educated should have larger capital gains rates on assets like housing, that is also biasing down their estimated consumption.

Table 5. Lifecycle Flows Per-Capita (\$2016), by Permanent Income and Age						
	Bottom 50 Percent					
	<35	35-44	45-54	55-64	65-74	75+
Income	18,597	30,257	30,714	26,811	25,493	26,286
Capital Gains	2,044	4,634	6,060	6,035	5,671	5,882
Net Interfamily Transfers Received	822	1,110	1,957	3,328	4,220	5,362
Net Change in Wealth	4,538	6,137	7,848	9,944	3,437	(505)
Saving	1,672	394	(169)	580	(6,454)	(11,749)
Consumption	16,926	29,863	30,883	26,231	31,947	37,957
	50th to 90th Percentiles					
	<35	35-44	45-54	55-64	65-74	75+
Income	33,914	58,081	64,950	67,800	54,020	49,054
Capital Gains	5,484	13,712	17,081	20,647	18,249	15,047
Net Interfamily Transfers Received	2,251	2,689	3,717	6,343	10,046	11,034
Net Change in Wealth	14,665	29,824	36,645	31,272	16,181	8,521
Saving	6,930	13,423	15,847	4,281	(12,115)	(17,560)
Consumption	26,984	44,658	49,103	63,519	66,135	66,527
	Top 10 Percent					
	<35	35-44	45-54	55-64	65-74	75+
Income	70,734	184,868	258,571	301,899	298,693	203,581
Capital Gains	36,259	95,983	135,220	148,607	127,670	122,752
Net Interfamily Transfers Received	4,315	9,545	18,628	36,172	54,641	55,718
Net Change in Wealth	84,187	172,162	232,037	202,656	107,596	121,498
Saving	43,613	66,634	78,189	17,877	(74,715)	(56,972)
Consumption	27,121	118,234	180,382	284,022	373,408	258,539

Lifecycle Flows by Permanent Income

The SCF oversampling strategy makes it possible to achieve a fair amount of statistical precision at the top of the permanent income distribution. Table 5 separates each birth cohort into three permanent income groups: the bottom 50 percent 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 that are largely offset by associated debts. The 50th through 90th percentiles of the permanent income distribution has substantial wealth holding, especially as they accumulate retirement balances and pay down housing and vehicle debt. The top 10 percent has incomes that far exceed the other permanent income groups, and they hold an even more disproportionate share of wealth, especially owned businesses.

Differences across agent-types by permanent income are clearly starker than differences by education, but most of the general observations about lifecycle patterns are the same. Income gaps by permanent income are large at young ages, and the differences in lifecycle income

profiles are now even starker, with the top 10 percent showing substantially more income growth over the lifecycle. Still, the broad lifecycle patterns of wealth change, conventionally defined saving, and interfamily transfers are the same across the three income groups. Most notably, there is substantial dissaving at older ages, with capital gains and net interfamily transfers accounting for continued wealth growth among retirees at all incomes.

The observed differences in lifecycle behavior by education (table 4) are even starker in the disaggregation by permanent income (table 5). Conventionally defined saving in the top 10 percent of the permanent income distribution (even as a share of disposable income) is far greater than for the other 90 percent, and that leads naturally to much higher capital gains relative to income later in the lifecycle. For the reasons mentioned above, some of the estimated saving of the young high income is likely due to the bias from missing *inter vivos* transfers in the highest permanent income group. That suggests the consumption of the young top 10 percent is biased down, offset by an upward bias in consumption for the top 10 percent later in the lifecycle. However, even with adjustments within the top 10 percent (shifting of consumption from old to young, offset by shifting saving from young to old) the estimated consumption at various lifecycle stages for the top 10 percent is far greater than what one finds in household spending data, and suggests that consumption inequality is on par with income inequality.

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 separate 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 conventionally defined savings playing a large positive role for some agent-types early in life, and is negative at older ages.

The results here reconcile apparent anomalies in prior empirical work on lifecycle saving and consumption, and provide important insights about whether and how particular modeling strategies are able to capture determinants of wealth inequality. Our pseudo-panel approach

makes it clear why wealth does not decline at older ages, even though conventionally defined saving shows a clear lifecycle pattern, including substantial drawdown of accumulated assets. Capital gains and interfamily transfers dominate wealth accumulation, which means that an empirical strategy of using measured labor income as the key source of funds for wealth accumulation and modeling a standard consumption/saving tradeoff will fall short of generating realistic wealth inequality in a Bewley-type model. The interpretation and modeling of differential capital gains—especially on owned businesses—is clearly key for modeling wealth inequality. Whether such gains are modeled as heterogeneous returns or untaxed and mislabeled labor income is another question, but in some sense, second order. The key insight is that the flow of income that model agents are presumed to allocate between saving and consumption is simply not being captured in the traditional income-oriented data sets used to calibrate models.

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 help guide this research going forward. The narrow focus on marketable wealth in this and previous work suggests that any distributional analysis based on those concepts will be missing the impact of important social insurance policies at lifecycle frequencies.⁴⁰ 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 (for example) measuring the change in net present value of social insurance on an accrual basis could fundamentally change inferences about overall saving behavior.⁴¹ Finally, although it is beyond the scope of this paper, the other possible avenue for productive future research is to use the pseudo-panel approach here to study consumption and saving behavior at business cycle frequencies. The SCF is a relative small and triennial data set, but combining the SCF with other higher-frequency data and directly addressing the issue of sampling variability through population smoothing or other techniques will make it possible to sort out consumption and saving responses to income and wealth shocks across various types of agents.

⁴⁰ The current state of market-based distributional analysis is well captured in the DINA approach of Piketty et al. (2018) or EG DNA approach of Zwijsenburg (2019). 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.

⁴¹ 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)
- Auclert, Adrian and Matthew Rognlie. 2018. "Inequality and Aggregate Demand." Working Paper No. 24280, Cambridge, MA: National Bureau of Economic Research.
- Avery, Robert B., Daniel J. Grodzicki and Kevin B. Moore. 2015. "Death and Taxes: An Evaluation of the Impact of Prospective Policies for Taxing Wealth at the Time of Death," *National Tax Journal*, 68(3): 601-632. (September)
- Bach, Laurent, Laurent E. Calvet, 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.
- Batty, Michael, Jesse Bricker, Joseph Briggs, Elizabeth Holmquist, Susan McIntosh, Kevin Moore, Eric Nielsen, Sarah Reber, Molly Shatto, Kamila Sommer, Tom Sweeney, and Alice Henriques Volz. 2019. "Introducing the Distributional Financial Accounts of the United States," Federal Reserve Board: FEDS Working Paper 2019-017. (March)
- 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)
- Benhabib, J. A. Bisin, and Mi Luo. 2019. "Wealth Distribution and Social Mobility in the US: A Quantitative Approach," *American Economic Review*, 109(5): 1623–1647. (May)
- Bhandari, Anmol, Serdar Birinci, Ellen R. McGrattan, and Kurt See. 2019. "What Do Survey Data Tell Us about U.S. Businesses?" Federal Reserve Bank of Minneapolis Staff Report 568. (January)
- 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)
- 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.
- 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)
- 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.
- Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi. 2018. "Stock Market Returns and Consumption." Working Paper No. 24262, Cambridge, MA: National Bureau of Economic Research.
- 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.
- Fagereng, 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.
- 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.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption Over the Lifecycle," *Econometrica*, 70(1): 47-89. (January)
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK," *American Economic Review*, 108(3): 697–743. (March)
- Krusell, Per and Anthony A. Smith, Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy," *The Journal of Political Economy*, 106(5):867-896.
- Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2019. "Income and Wealth Inequality in America, 1949-2016," Working Paper, University of Bonn. (April)
- 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.
- 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):519–78.

Smith, Matthew, Danny Yagan, Owen M. Zidar, Eric Zwick. 2019. “Capitalists in the 21st Century,” Working Paper No. 25442, Cambridge, MA: National Bureau of Economic Research.

Zwijnenburg, Jorrit. 2019. “Unequal Distributions: EG DNA versus DINA Approach,” *American Economic Review*, 109: 296–301.

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.⁴²

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

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

⁴² 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 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 the FA measure is somewhat noisier.⁴³ 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

⁴³ See FRB Z1 release, Table F.6, lines 41 and 45, available at <https://www.federalreserve.gov/releases/z1/current>.

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.

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.⁴⁴ 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.⁴⁵ 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

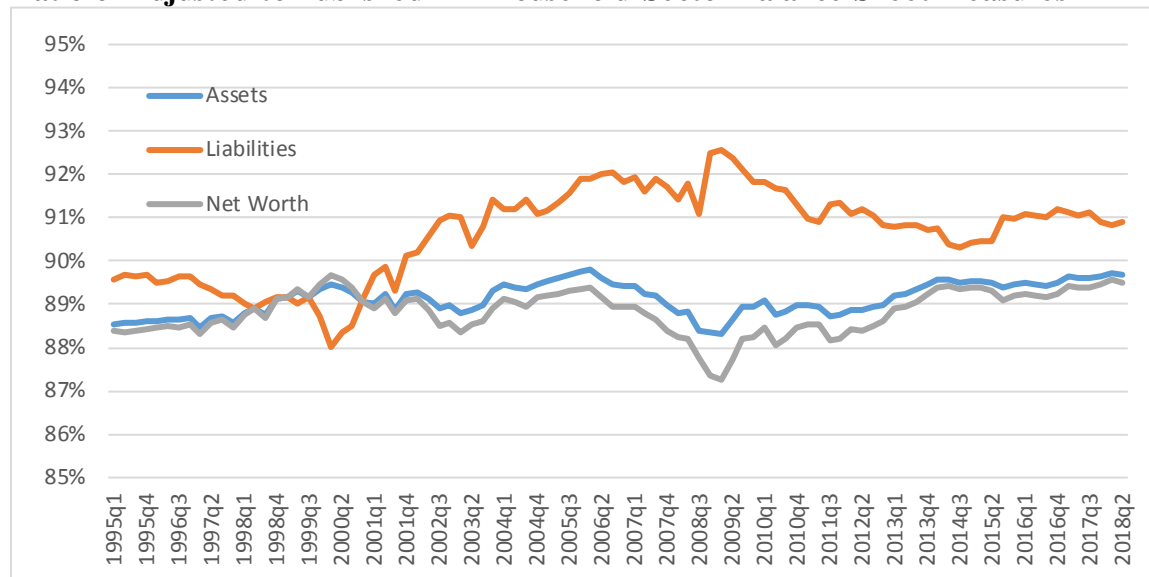
⁴⁴ See Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus. 2016. “Measuring Income and Wealth at the Top Using Administrative and Survey Data,” *Brookings Papers on Economic Activity*, 1:2016, p. 261-321, and Dettling, Lisa J., Sebastian Devlin-Foltz, Jacob Krimmel, Sarah Pack, and Jeff 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).

⁴⁵ All published FA series are accessed through the Federal Reserve Board web site at <https://www.federalreserve.gov/releases/z1/current>.

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.

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.⁴⁶

⁴⁶ Details on the SCF DB pension imputation are available in 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. 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 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 2016.⁴⁷ 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:

⁴⁷ For an introduction to the SCF and overview of 2016 results, see Jesse Bricker, Lisa J. Detling, 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.⁴⁸ 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.⁴⁹ 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

⁴⁸ 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.

⁴⁹ 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.⁵⁰ 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.⁵¹

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

⁵⁰ 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.

⁵¹ 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).⁵² 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.⁵³

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.⁵⁴

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.

⁵² The CBO differential mortality model is described in Working Paper 2007-11, August 2007, available at <https://www.cbo.gov/publication/19096>.

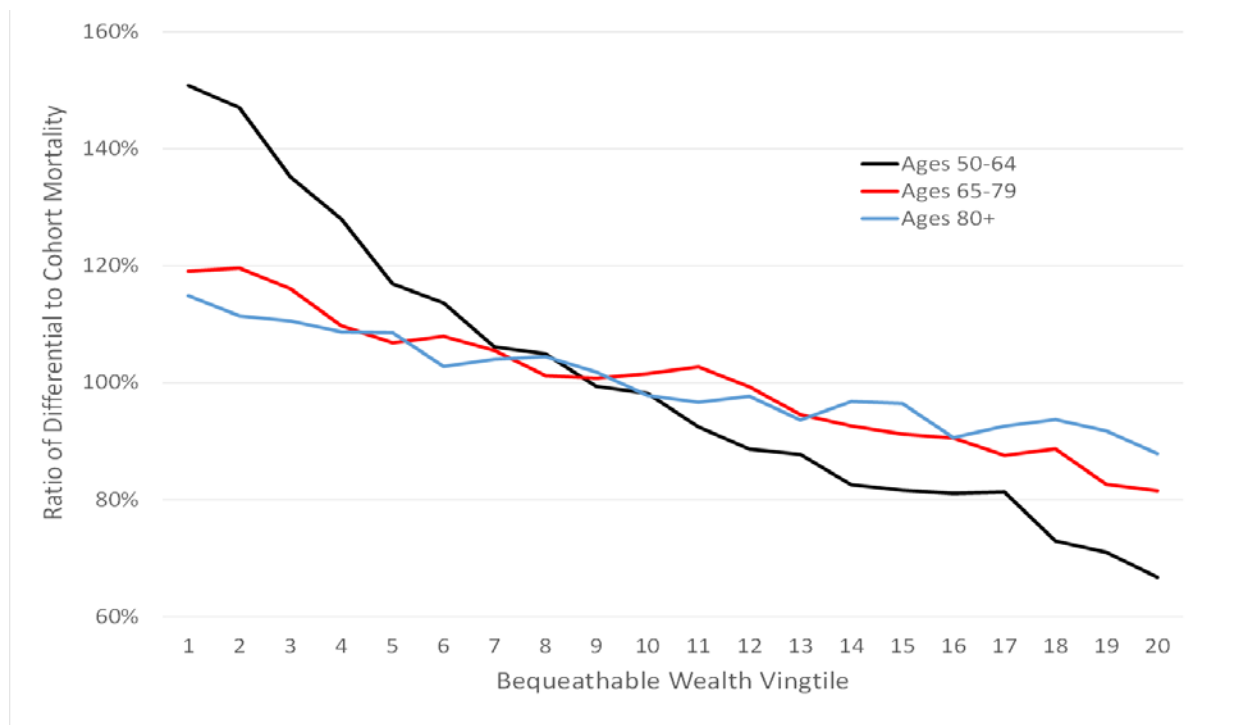
⁵³ The only notable difference between socio-economic measures is in the construction of the income variable. CBO's estimates are based on SSA earnings and death records linked to various Survey of Income and Program Participation (SIPP) data files. The income quintile variables used by CBO are based (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.

⁵⁴ 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.⁵⁵ 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

⁵⁵ 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.

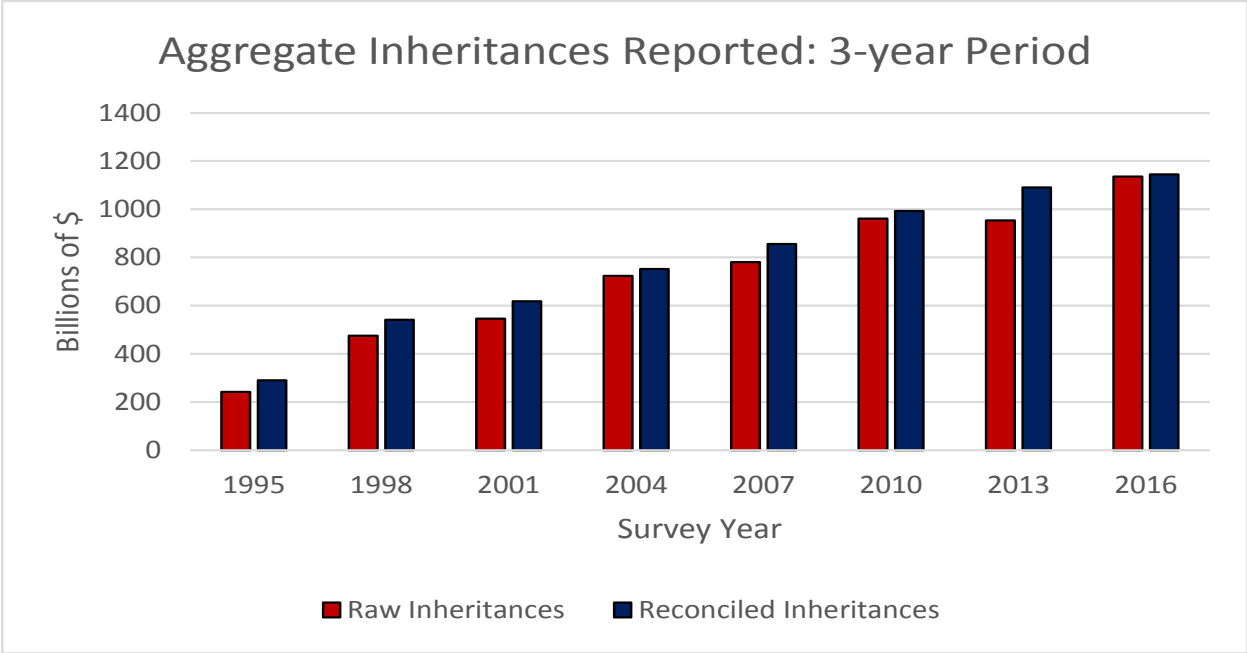
Inheritances of non-property (financial assets and cash) are included in the SCF income concept if received in the year prior to the survey, which is the same time frame for measuring the other income flows (and, we will see, most *inter vivos* transfers). The income section actually occurs before the 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

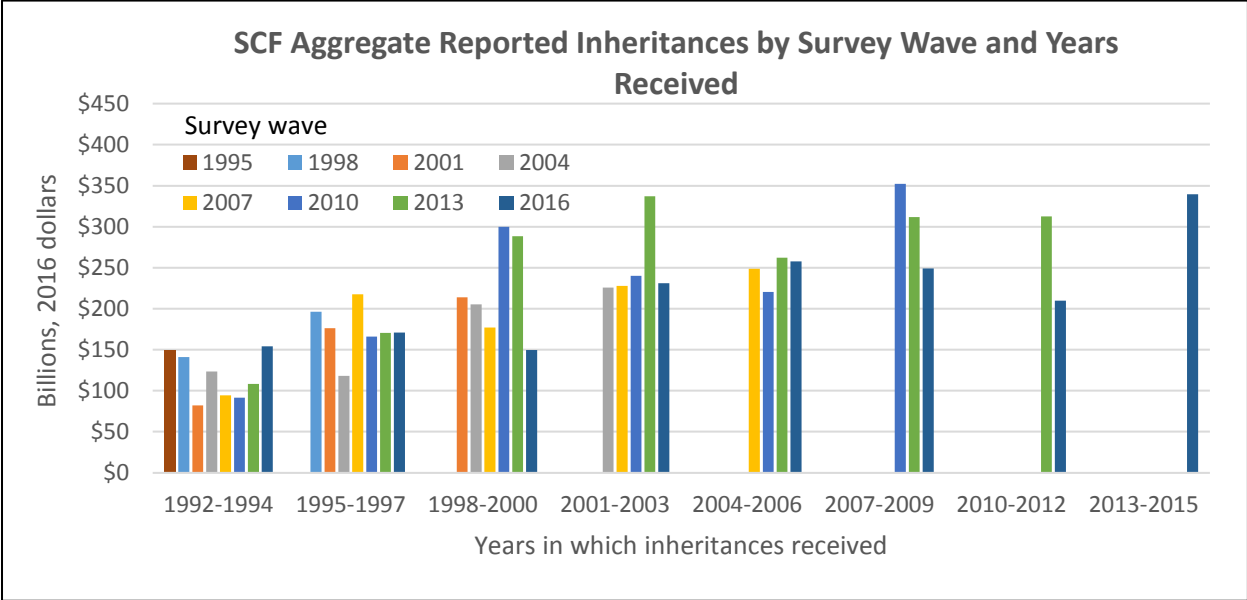
In practice, the primary inheritance questions do a good job of capturing all inheritances and gifts, and the adjustments we made for incremental information from the assets and income modules add relatively little to the total inheritance estimates. This observation in large part reflects the key SCF 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.⁵⁶ 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.

⁵⁶ 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.⁵⁷

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.

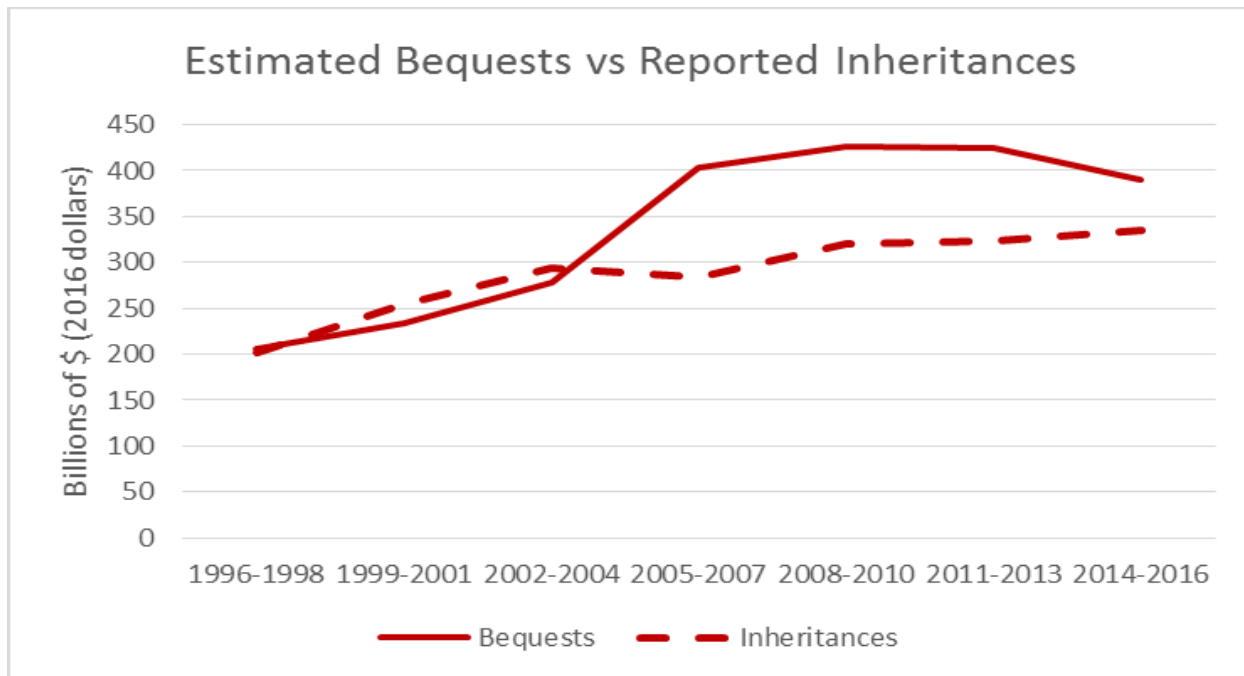
Comparing Estimated Bequests Made and Reported Inheritances Received

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-

⁵⁷ 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.⁵⁸ 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.

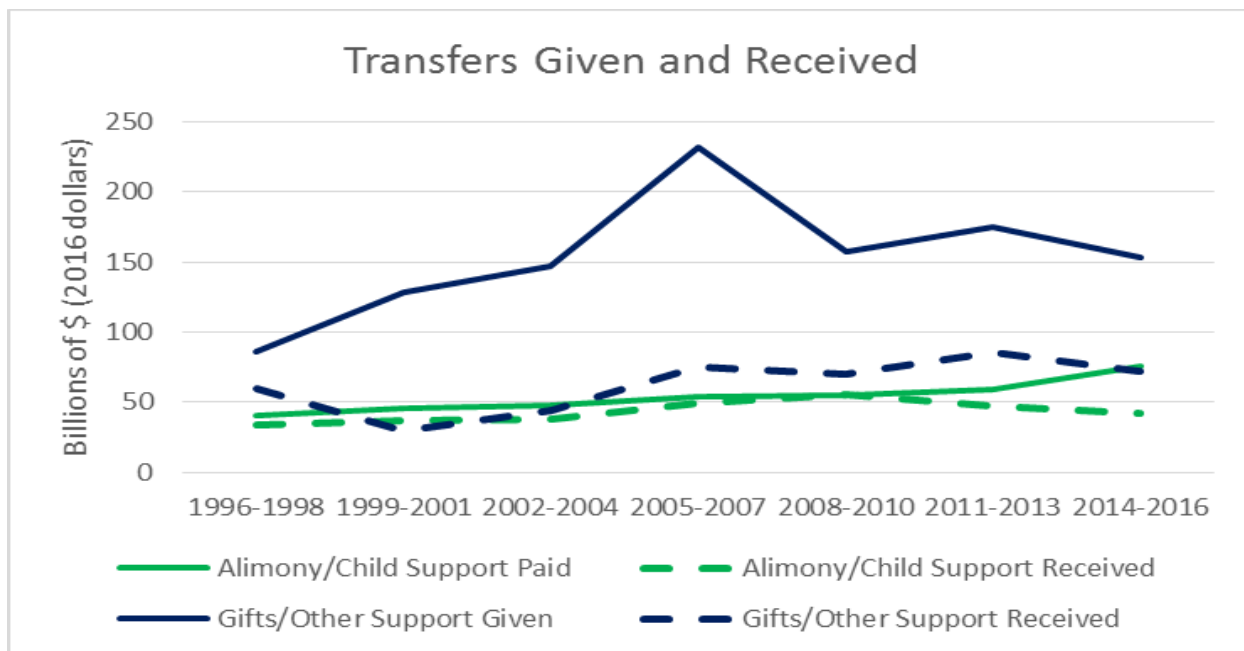
Reconciling Inter Vivos Transfers Made and Received

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:

⁵⁸ 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?