FISCAL STUDIES, vol. 37, no. 1, pp. 153–192 (2016) 0143-5671

Heterogeneity in Economic Shocks and Household Spending in the US*

SEBASTIAN DEVLIN-FOLTZ† and JOHN SABELHAUS‡

†*Board of Governors of the Federal Reserve System (Sebastian.J.Devlin-Foltz@frb.gov)*
‡*Board of Governors of the Federal Reserve System (john.sabelhaus@frb.gov)*

Abstract

Large swings in aggregate household sector spending, especially for big-ticket items such as cars and housing, have been a dominant feature of the macroeconomic landscape in the past two decades. Income and wealth inequality increased over the same period, leading some to suggest the two phenomena are interconnected. Indeed, there is supporting evidence for the idea that heterogeneity in economic shocks and spending are connected, most notably in studies using local-area geography as the unit of analysis. The Survey of Consumer Finances (SCF) provides a household-level perspective on changes in wealth, income and spending across different types of families. The SCF confirms that inequality is indeed increasing in recent decades, and the data provide support for the proposition that shocks to income and wealth are indeed related to large swings in spending across and within birth cohorts. However, the economic shocks associated with the Great Recession and changes in

*Submitted September 2015.

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. The authors are grateful for comments from Thomas Crossley, Cormac O'Dea and an anonymous referee, for the lively discussions with participants at Brookings Institution and Federal Reserve Board brown bag presentations, and for feedback received during the March 2015 conference on 'Household Wealth Data and Public Policy', organised by the Institute for Fiscal Studies and Public Economics UK, sponsored by the Nuffield Foundation and hosted by the Bank of England.

Keywords: inequality, consumption, deleveraging, synthetic panels.

JEL classification numbers: C8, D1, E2.

Fiscal Studies © 2016 Institute for Fiscal Studies. Published by John Wiley & Sons Ltd, 9600 Garsington Road, Oxford, OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA.

spending and debt to income ratios are widespread, and inconsistent with a narrow focus on the experiences and changes in behaviour of particular (especially low- and modest-income) households.

Policy points

- Shocks to income and wealth, and corresponding changes in spending behaviour, were widespread across the population in the period leading up to, during and after the Great Recession.
- The boom and bust in household borrowing and spending was not simply driven by policies expanding access to homeownership and credit for previously under-served groups.
- When beliefs about house prices and other economic fundamentals were proven wrong, and widespread behavioural changes ensued, the families who became unable to meet their financial obligations experienced the brunt of the collateral damage.

I. Introduction

One of the continuing legacies of the Great Recession is a dramatic slowdown in the growth rate of aggregate consumer spending. The ongoing slowdown in consumption growth follows a much larger decrease in spending than had occurred in other recent US recessions, most notably for big-ticket items such as cars and owned housing. The dramatic decline in spending during the Great Recession, in turn, had followed a fairly dramatic surge in consumption, housing investment and household debt during the decade or so preceding the financial crisis. This boom and bust in aggregate household spending has occurred at the same time that income and wealth inequality are rising, leading some to conclude that the phenomena are interconnected. This paper uses household-level data on income, wealth, debt and spending to investigate the connection between inequality and macroeconomic outcomes.

The idea that rising inequality and the observed swings in aggregate spending are related has found support in both theory and empirical research. Theory suggests several channels by which rising inequality may have interacted with underlying macroeconomic trends to generate the boom and bust in spending. For example, some low- to modest-income and/or credit-constrained families may have reacted to their relatively slow income growth during the boom by borrowing more, which was made possible by rising house prices, changes in lending standards and a macroeconomic environment characterised by low unemployment and stability. Those same families then reacted to the collapse in housing prices by dramatically reining in their spending, either by choice or because they were constrained from obtaining credit in the post-crisis economy.

Fiscal Studies © 2016 Institute for Fiscal Studies

The potential connection between rising inequality and spending fluctuations also has empirical support. Much of this evidence comes from studies based on geographically grouped data. For example, the boom and bust in consumption was most pronounced in lower-income areas where house prices were increasing, while borrowing and consumption in higher-income areas did not seem to respond to house prices.[1] Also, consumption during the boom grew the most in areas where inequality was rising fastest, suggesting some sort of 'keeping up with the Joneses' or 'status goods' effects might be affecting spending.[2]

In this paper, we use repeated cross-sections from the triennial Survey of Consumer Finances (SCF) to study income and wealth shocks along with borrowing and spending responses at the household level.[3] The SCF data are cross-sections, and thus we construct synthetic panels by age and permanent income / education in order to directly address the inequality and growth narratives. Specifically, we construct life-cycle trajectories for income, wealth, borrowing and big-ticket (cars and owned housing) spending measures. This makes it possible to identify how various birth cohorts, and income or education groups within those birth cohorts, were affected by the boom and bust in income and wealth, and how they responded in terms of borrowing and spending.

The constructed life-cycle trajectories provide some support for certain aspects of the inequality narratives, but the overall impression is that more widespread shocks to income and wealth, along with fundamental changes in behaviour, have taken place across all permanent income groups. The supporting evidence for the inequality narrative comes from the fact that the *relative* shocks to (especially) wealth at the bottom of the income distribution were larger, and *relative* responses (car buying and owned-home transactions) were also larger for those same groups. However, the overall narrative about inequality, shocks and spending responses is not consistent with the fact that the behavioural changes across income groups were widespread, and the differential responses of the lower-income families account for a very small share of the overall change in (for example) new car buying.

To a large extent, this 'source of change' decomposition analysis by income groups is driven by the fact that (for example) in 2007, the bottom half of families sorted by our measure of permanent income accounted for only 19 per cent of new car spending. The fact that the bottom half of families then accounts for 36 per cent of the *decline* in new car spending between 2007 and 2010 suggests that a differential response for the bottom half of families did take place, but the top half of the permanent income distribution still accounts for the other 64 per cent of the dramatic decline in new car buying.

[1]Mian and Sufi, 2014a.
[2]Bertrand and Morse, 2013; Bricker, Ramcharan and Krimmel, 2014.
[3]For an overview of the SCF and latest results, see Bricker et al. (2014).

In the various decompositions presented here, our focus is consistently drawn to shocks and spending responses among middle-aged and middle- to high-income families, who account for the lion's share of economic activity. Those groups experienced large income and wealth shocks, and they responded as expected.

Trends in homeownership, owned-housing transactions and household debt also provide mixed evidence about the inequality narrative, but the data again suggest more widespread behavioural changes are at work. Movement into homeownership and debt growth during the boom period changed noticeably (from a life-cycle perspective) for the bottom half of families during the boom period (2001 to 2007), but the top half of the distribution also saw dramatic increases in owned-housing transactions and debt growth, especially among families above median income but below the top 5 per cent (what we call the 'next 45' permanent income group). The overall growth in household debt between 2001 and 2007 was widespread, as evidenced by the fact that the shares of household debt outstanding for the various permanent income groups in 2007 were nearly identical to the shares of debt outstanding in 2001. The ratios of debt *to income* at the very top did not grow as much between 2001 and 2007, which does provide some support for the inequality narrative, because (in a sense) one can argue that the very-highest-income families took on (*ex post*) manageable debt, while the other 95 per cent were (*ex post*) borrowing beyond their means.

The post-2010 observations on spending and owned-housing activity reinforce the idea that widespread and fundamental changes in spending behaviour have taken place. The dramatic decline in car buying between 2007 and 2010 was followed by a modest recovery between 2010 and 2013, but the slowdown in spending (at least relative to pre-recession levels) continued to be widespread. Likewise, owned-housing turnover declined across all income groups and has remained at greatly diminished levels. These widespread changes in behaviour are consistent with economic fundamentals such as diminished expectations about permanent income and/or future house prices, or possibly increased uncertainty about those expectations. The changes in behaviour associated with the inequality narrative are certainly complementary to changes in expectations and other economic fundamentals, but the inequality narrative by itself seems to fall well short of explaining recent macroeconomic fluctuations.

II. The link between inequality and macroeconomic fluctuations

There are a number of explanations for the dramatic swings in aggregate household spending over the past two decades, and the evolving distributions of income and wealth play an important role in many of those narratives. The macro fundamentals are by now very well known. In the decade or so

Fiscal Studies © 2016 Institute for Fiscal Studies

leading up to the Great Recession, asset prices – especially for real estate – rose dramatically. Also during the boom, spending on both housing and non-housing goods and services rose dramatically, and household debt – again, most notably for housing – rose much faster than income. When house prices began to drop after 2006, household spending dropped precipitously, especially for durable goods and housing, and the recession officially began in late 2008. Now, several years after the official end of the recession, spending continues to grow much more slowly than in the pre-recession period, and the slow growth of consumption has led to a reduction in aggregate household debt.

The possibility that rising inequality may be fuelling or at least amplifying these macroeconomic fluctuations has some support in both theory and empirical research. For example, Kumhof, Rancière and Winant (2015) provide some basic empirical observations connecting inequality, the distribution of debt and economic crises.[4] They then go on to explain the connection between those phenomena using a parsimonious model in which higher-income families have more of a taste for wealth than low- and middle-income families.[5] That differential taste for wealth leads higher-income families to increase lending to non-wealthy families when permanent income shocks increase the gap between rich and poor. The increased debt among low- and middle-income families, and associated exposure to economic shocks, set in motion the boom and bust cycle that characterises financial crises.

The question of whether this type of inequality-driven borrowing and spending mechanism is underlying recent trends is debatable, however. For example, some papers have argued that the observed dramatic swings in aggregate household spending can be reconciled with relatively simple representative-agent versions of the life-cycle/permanent-income model. De Nardi, French and Benson (2012) use a representative-agent framework and argue that, for reasonable reduced-form parameter values, the shocks to asset values alone can explain about one-third of the departure of aggregate consumption from trend. In addition, diminished levels of expected permanent income can explain the remainder (if not more) of the drop in consumption, though that conclusion is somewhat sensitive to the time horizon for income expectations.

Other research does point towards an important distributional component, though the specific mechanism and even the direction of change for various income and wealth groups are not (at least on the surface) consistent with the simple connection between inequality and borrowing described above. For example, Petev, Pistaferri and Eksten (2012) show that consumer confidence

[4]Cynamon and Fazzari (2014) come to a similar conclusion about debt and spending across the income distribution, using a combination of micro and macro data sources.

[5]Bertrand and Morse (2013) and Bricker, Ramcharan and Krimmel (2014) describe other (non-standard) mechanisms by which rising inequality may increase aggregate spending.

and consumption fell disproportionately more for higher-income families, confirming the idea that the wealth channel was probably a key for them, and their decreased spending is a key to understanding consumption dynamics in the Great Recession. But the authors note that spending also fell for consumers with little or no balance-sheet wealth, suggesting some other factors were also important – possibly disproportionately diminished income expectations, high propensities to consume out of transitory income for some groups, or credit constraints.

Christelis, Georgarakos and Jappelli (2015) likewise stress the importance of wealth effects on consumption, especially among those consumers who perceived that losses on financial assets would be permanent, but they also find large unemployment effects on spending, again suggesting larger-than-expected consumption responses to transitory income shocks for low- and middle-income families. Jappelli and Pistaferri (2014) confirm that there is enormous heterogeneity in consumption responses to (hypothetical) transitory income shocks, and the estimated responses are particularly large for those with low levels of cash-on-hand. Some heterogeneity by levels of cash-on-hand is predicted by standard models with uncertainty about future income, but the authors conclude that additional departures from standard models (very high discount rates or myopia for low-wealth consumers) are required to fully explain the range of marginal consumption responses.

It is also worth noting that *observed* binding credit constraints for low-wealth households are not essential to generate substantial spending reactions to transitory income fluctuations. Crossley and Low (2014) combine data on actual labour market experiences with self-reported income expectations, and show that the potential for *future* credit constraints is also an important driver of spending behaviour. Families reporting binding constraints do react more to income shocks, but the mechanism by which more widespread consumption responses could occur is confirmed by the self-reports. Another theoretical exception to the idea that fluctuations are driven only by (observably) credit-constrained families overreacting comes out of two-good models with high transaction costs on one of the goods. Chetty and Szeidl (2007) provide the theoretical basis for 'wealthy hand-to-mouth' behaviour, and Kaplan and Violante (2014) and Kaplan, Violante and Weidner (2014) provide empirical support for the proposition that families with (observably) high wealth will also change spending quite dramatically when transitory income shocks occur.[6]

These various inequality narratives share some general predictions about the distribution of income, borrowing and spending over the past two decades, even though the specific mechanisms underlying the stories are different. The idea is that we should have seen borrowing and spending increases during the boom moving in a highly correlated way across groups, as the types of families

[6]Berger and Vavra (2015) present a similar argument based on the costs of adjusting consumer durables.

Fiscal Studies © 2016 Institute for Fiscal Studies

who previously were unable to obtain credit found they could do so during the housing boom and associated period of credit liberalisation. Those same families then saw their access to credit greatly restricted after housing prices collapsed, and the collapse of their spending caused both the recession and continued slow growth.

The most compelling evidence that the changes in borrowing and spending were concentrated among certain population subgroups has been put forth by Mian and Sufi (2014a). They use geographically-constructed data to show that spending out of rising home equity during the boom was most pronounced in areas where income is low and house prices increased, suggesting that previously credit-constrained consumers led the spending boom. Conversely, borrowing and spending in high-income areas were unresponsive to changes in house prices. This 'housing net worth channel' then permeates through the rest of the economy, with the largest effects in non-tradable sectors, as shown by Mian and Sufi (2014b). One challenge to the simple distributional story comes from Justiniano, Primiceri and Tambalotti (2015), who agree with the collateral channel effect, but show the patterns of debt to income are consistent with a model with increased demand for housing, not simply an exogenous relaxation of credit constraints. Thus, the question to be addressed using the SCF is whether we see the patterns of debt and spending by income suggested in the various inequality-based narratives, versus something more widespread in nature.

III. Tracking household sector aggregates using the SCF

The first step in using the Survey of Consumer Finances for studying fluctuations in income, wealth, spending and debt at a micro level is verifying that the survey does indeed capture the macroeconomic phenomena of interest.[7] The specific goal of this section is to show that the SCF tracks aggregate household sector incomes, household net worth, spending on new cars, recent home buying and household debt. The focus is on the period leading up to, during and following the US financial crisis, 1995–2013. In general, the results are very promising, and suggestive that the SCF sampling and survey strategy captures macroeconomic trends and fluctuations quite well over the period being studied.[8]

[7]Dettling et al. (2015) provide a much more thorough comparison of the various household sector aggregates in the SCF versus other administrative micro and macro data sources. In many ways, this section of our paper is a very high-level treatment of that more extensive reconciliation exercise.

[8]Bricker et al. (2015) show that the SCF also tracks changes in the distribution of income and wealth observed in other data sets, such as those analysed by Piketty and Saez (2003 and 2015) and Saez and Zucman (2014) over the same period, after adjusting for conceptual differences.

1. Income

The concept of income in household surveys is not the same as the concept of income being measured in the National Income and Product Accounts (NIPA) or even the concept in other micro administrative data, such as those derived from Internal Revenue Service (IRS) Statistics of Income (SOI) tax records. However, after conceptually adjusting to the extent possible, the SCF generally tracks NIPA and SOI aggregates (Figure 1).[9] Over the five survey waves between 1995 and 2007, aggregate SCF income slightly more than doubled, while the NIPA and SOI aggregates almost exactly doubled. Between the 2007 and 2013 surveys, SCF income grew by only about 10 per cent, while the NIPA total grew roughly 20 per cent and SOI incomes grew 13 per cent. Most of the divergence over the six-year period occurred in the first half, as both SCF and SOI income were lower in the 2010 survey than in the 2007 survey.

In addition to the markers for actual reported aggregate SCF incomes (the diamonds) on Figure 1, there is a second set of markers for SCF 'usual' income (the circles). The concept of usual income is a respondent self-reported measure of 'permanent' income, and thus abstracts from transitory fluctuations.[10] This measure is crucial to our within-cohort decomposition strategy below, and we

FIGURE 1

Aggregate income

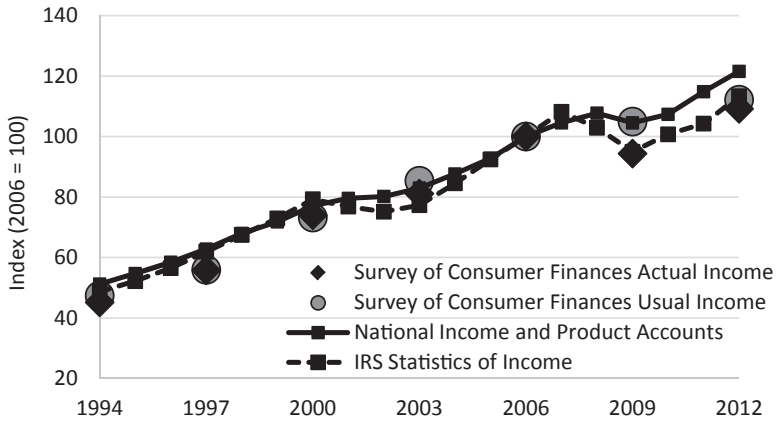

Source: National Income and Product Accounts; Statistics of Income; Survey of Consumer Finances. See the online Data Appendix for descriptions.

[9]SCF incomes are measured for the year prior to the triennial survey. Therefore Figure 1 shows SCF aggregates in 1994, 1997, etc. through 2012. For details about the specific conceptual adjustments being applied to both the SCF and aggregate data sources, see the online Data Appendix and Dettling et al. (2015).

[10]See Sabelhaus and Ackerman (2012) for a detailed analysis of the usual income measure and its usefulness for understanding fluctuations in spending.

Fiscal Studies © 2016 Institute for Fiscal Studies

discuss how the measure is constructed and its statistical properties in Section IV.2. For now, it is worth noting that, in general, the actual and permanent income *aggregates* in the SCF track each other across surveys, *except* in the 2010 survey when negative (self-reported) transitory shocks were much more common than positive shocks. In other survey years, the negative and positive shocks across families largely cancelled out, and the aggregate measures were basically the same.

2. Net worth

The aggregate benchmark for evaluating how well SCF aggregates are capturing recent trends in net worth is the Financial Accounts of the United States (FA).[11] As with aggregate incomes, the SCF does a good job tracking overall trends in household sector net worth, after adjusting to the extent possible for conceptual differences between SCF and FA (Figure 2). The substantial and generally sustained growth in household sector net worth between 1995 and 2007 is reflected in both data sets, with wealth more than doubling over the period. The SCF captures the two phases of rising household sector wealth, as stock prices fuelled the boom between 1995 and 2001 and housing prices fuelled the boom between 2001 and 2007.

FIGURE 2

Aggregate net worth

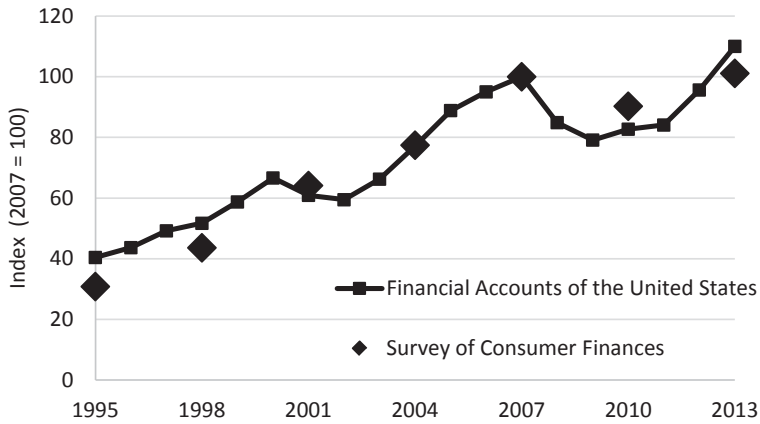

Source: Financial Accounts of the United States; Survey of Consumer Finances. See the online Data Appendix for descriptions.

[11]The Financial Accounts, produced quarterly by the Federal Reserve Board, were formerly known as the Flow of Funds Accounts.

Timing and lagged valuations are likely explanations for most of the relatively modest divergence between net worth measures after 2007. The SCF wealth numbers did not fall as much as FA between 2007 and 2010, but then did not increase as much as FA between 2010 and 2013. Some of this is attributable to respondent recall, because survey participants are asked to estimate values for assets that are changing (sometimes rapidly) in value, such as housing and corporate equities. If the respondents are using what is likely outdated information, meaning the last time they checked an account or observed a real-estate transaction in their neighbourhood, they will tend to understate losses in asset price downturns (such as 2007–10) and understate gains in asset price recoveries (2010–13).

3. New car spending

Spending on new cars is a relatively straightforward concept, but inferences about the exact timing of new car purchases in the SCF are made indirectly, so the time series is not completely comparable to published aggregates from the NIPA. SCF car buying is inferred based on measuring the stock of cars and asking about model year and (if a loan exists) when the loan was taken out. In addition, the separation between cars for personal use and cars for business use (which are excluded from the SCF) confounds the comparison. Still, the long-run trends and dramatic fluctuations in new car spending are well captured in the SCF (Figure 3). Both the near-doubling of car spending between 1995 and

FIGURE 3

Aggregate new car spending

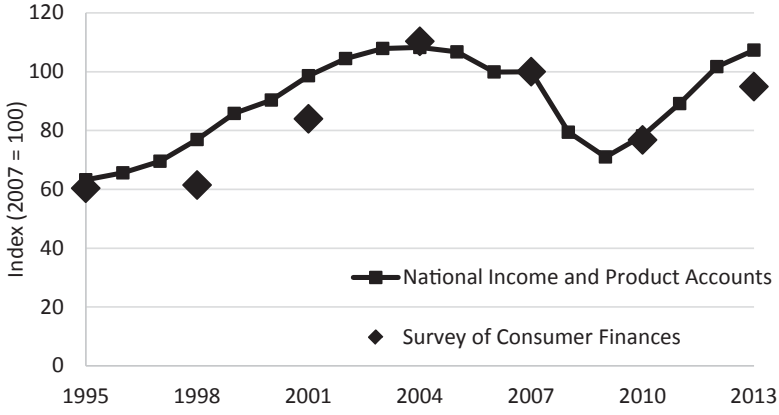

Source: National Income and Product Accounts; Survey of Consumer Finances. See the online Data Appendix for descriptions.

Fiscal Studies © 2016 Institute for Fiscal Studies

2004 and the implosion of car buying between 2007 and 2010 show up in the two data series.

4. Recent housing purchases

The concept of housing sector activity captured in the SCF has no direct analogue to a NIPA measure of spending or investment, but it is nonetheless a useful indicator of the pace of economic activity. In the SCF, respondents are asked whether they own their home and, if so, how long they have lived there. That provides a measure of housing turnover, not housing investment or rental value per se. But, as corroborated by comparable housing turnover measures from the National Association of Realtors and the Department of Commerce, the SCF is capturing the boom and bust in housing transactions over the past two decades (Figure 4). Indeed, the pace of housing turnover is highly correlated with gross residential investment in the NIPA.

5. Household debt

Given the importance of household debt in narratives about inequality and spending over the past two decades, the final aggregate comparison focuses on the debt component of household sector net worth. Household debt more than tripled in both the SCF and the conceptually comparable FA data series between 1995 and 2007, with most of the growth occurring between 2001 and 2007 (Figure 5). The SCF also tracks the (modest) deleveraging that occurred

FIGURE 4

Aggregate recent housing purchases

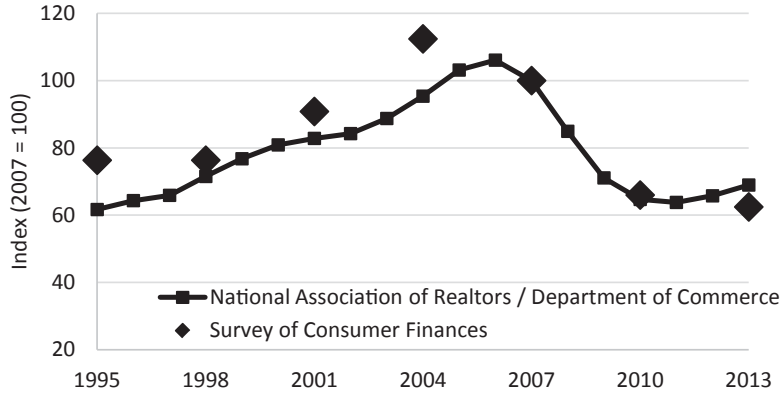

Source: National Association of Realtors; Department of Commerce; Survey of Consumer Finances. See the online Data Appendix for descriptions.

FIGURE 5
Aggregate household debt

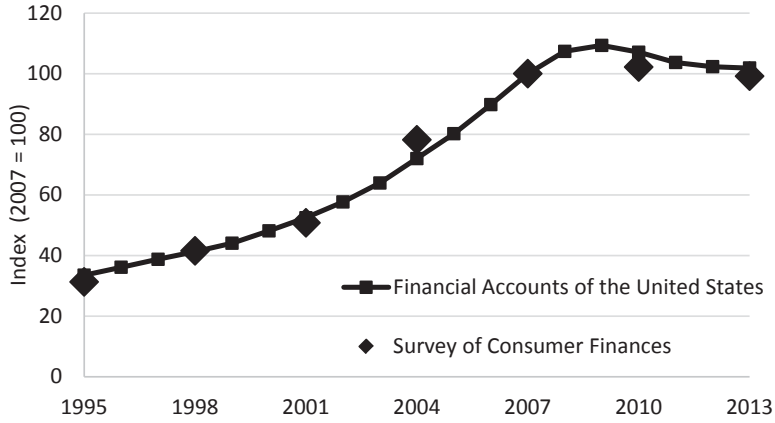

Source: Financial Accounts of the United States; Survey of Consumer Finances. See the online Data Appendix for descriptions.

in debt between 2007 and 2013. Debt has roughly stabilised, and debt burdens (relative to income) moved down in the two most recent surveys.

IV. Constructing synthetic panels using the SCF

The Survey of Consumer Finances is a repeated cross-section, and thus decomposing trends and fluctuations in income, wealth and spending involves creating a synthetic panel and then looking at outcomes across groups and time.[12] The synthetic panel approach is based on the principle that one does not need to track the same families over time, because changes in group-level statistics (means, medians or other fractiles) provide the desired information about differences across groups. The most obvious top level of disaggregation for our purposes is birth cohort, because each birth-cohort group systematically ages three years between SCF surveys, and changes in behaviour at various points in the life cycle are key for the inequality narratives. Within birth-cohort groups, we also decompose group-level changes by the SCF 'usual' income measure, which closely tracks permanent income, and by educational attainment. Thus we are able to look within birth cohorts across the income dimension.

[12]Deaton (1985) provides an excellent discussion of the issues involved with constructing synthetic panels and Moffitt (1993) adds a more dynamic perspective. For a closely-related and recent application of the synthetic panel approach to studying household sector finances, see Attanasio and Borella (2014).

Fiscal Studies © 2016 Institute for Fiscal Studies

1. Birth cohorts

We assign SCF households into birth-year cohort groups to track life-cycle developments in income, wealth and consumption. A household is assigned a cohort based on the reported birth year of the head of household in each survey year. The oldest birth cohort used in our analysis is comprised of households with heads born between 1931 and 1940, while the youngest cohort is comprised of households headed by people born between 1981 and 1990. There are four additional cohorts in between. We restrict our analysis to these six cohorts due to sample size issues for households headed by those born before 1930 or after 1990. In particular, with roughly 800 to 1,000 families per birth cohort in a given survey year, it is possible to create within-cohort groups based on income and education.

2. 'Usual' income groups

The usual income classifier is derived from survey questions about the gap between actual and 'normal' or 'usual' income in the SCF.[13] Towards the end of the SCF survey, after detailed income components have been summed to arrive at a total, respondents are asked whether that total income is 'unusually high or low compared to what you would expect in a "normal" year, or is it normal?'. Most respondents say their reported total income is in fact about normal – the median gap between actual and normal income is zero in every survey year. However, sizeable minorities of respondents indicate that their income is either unusually high or unusually low, and those fractions vary predictably and systematically with business-cycle conditions (Table 1). If the respondent indicates unusually high or low income, the survey has a follow-up question that asks for the level of usual income.

The canonical approach to deriving transitory income shocks involves using residuals of earnings or income equations estimated using panel data. Transitory shocks are solved for as one component of overall income change: the unexplained income change that does not appear (to the econometrician) to be permanent. Although the transitory income shocks in the SCF are estimated using a very different approach, the high-level statistical properties of the self-reported gaps between actual and normal income seem generally consistent with the properties of transitory income shocks derived from the residuals of estimated equations. In particular, the variances of the percentage gap between actual and normal income are of the same general magnitudes as the variances of residual-based annual transitory shocks, and the shape of the distribution of the gaps changes asymmetrically over the course of the business cycle in ways that are consistent with residual-based estimates. Thus, there is reason

[13]The SCF has maintained a consistent methodological design since the 1989 survey, though the question on 'normal' income was not added until the 1995 survey.

TABLE 1

Statistical properties of self-reported transitory income shocks

Incidence and mean and median shocks

Survey year	Families reporting actual income lower than normal income: Percentage of all households	Mean difference	Median difference	Families reporting actual income greater than normal income: Percentage of all households	Mean difference	Median difference
1995	17%	–$23,225	–$15,498	9%	$38,624	$15,544
1998	16%	–$26,162	–$15,751	10%	$68,838	$17,394
2001	14%	–$35,013	–$17,270	11%	$86,720	$20,251
2004	20%	–$30,754	–$15,198	9%	$57,866	$19,377
2007	14%	–$33,740	–$17,565	9%	$93,892	$17,551
2010	25%	–$38,575	–$18,302	6%	$64,633	$15,908
2013	18%	–$31,058	–$15,218	7%	$91,712	$15,328

Unconditional mean and variance

Survey year	All households: Average gap between actual and normal income	Households with actual and normal income > $0: Average gap between actual and normal income	Variance of percentage gap between actual and normal income[a]	Households with actual and normal income > $5,000: Average gap between actual and normal income	Variance of percentage gap between actual and normal income[a]
1995	–$614	–$308	0.127	–$149	0.083
1998	$2,691	$2,868	0.134	$2,988	0.091
2001	$4,368	$4,483	0.105	$4,610	0.085
2004	–$1,032	–$706	0.124	–$583	0.091
2007	$3,660	$3,910	0.108	$4,036	0.084
2010	–$5,812	–$4,980	0.145	–$4,836	0.109
2013	$1,063	$1,407	0.114	$1,540	0.093

[a]Variance of [ln(actual income)–ln(normal income)].

Note: All dollar values are shown in 2013 dollars.

Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

to believe that the households that self-report experiencing a transitory shock are the same households that the econometrician would identify as having experienced a transitory shock simply by looking at changes in their income over time.

Transitory income shocks have been estimated using various data sets, different income and earnings concepts, individual- and household-level units of observation, and alternative parameterisations of the stochastic process for the shocks themselves. A simple but descriptive specification involves decomposing log earnings or income (y_{it}) into a deterministic component that evolves with observable characteristics (x_{it}), a permanent component that evolves slowly over time (μ_{it}) and a transitory component (ε_{it}). That is,

$$(1) \qquad y_{it} = \beta x_{it} + \mu_{it} + \varepsilon_{it}.$$

The permanent component changes when the individual receives a permanent shock (η_{it}):

$$(2) \qquad \mu_{it} = \mu_{it-1} + \eta_{it}.$$

Given simplifying iid (independently and identically distributed) assumptions on ε_{it} and η_{it}, it is straightforward to recover estimates of the variances for the two shocks (σ_{ε}^2 and σ_{η}^2 respectively) using panel data.[14,15] Although there is a great deal of heterogeneity in underlying income concepts, unit of observation, data sources and methods, there is a fair amount of uniformity in the literature in estimates for the percentage variance of transitory shocks, with values generally below but near 10 per cent.[16]

[14]The essence of the method for separating permanent and transitory shocks, described succinctly in Carroll (1992), is to measure the variance of income changes at multiple frequencies, then acknowledge that every one of those variances has two transitory shocks (for each of the two years at the endpoints) and a number of permanent shocks equal to the frequency over which the change is being measured. Thus, the variance of one-year income changes has two σ_{ε}^2 terms and one σ_{η}^2, the variance of two-year income changes has two of each, the variance of income changes over three years has two σ_{ε}^2 and three σ_{η}^2 terms, etc. Given panel data with more than two years of data, one measures the variance of income change at every frequency then solves the (generally over-identified) system of equations for σ_{ε}^2 and σ_{η}^2. Although studies of income volatility often use more complex stochastic processes that allow transitory shocks to have effects that last more than a year, all of the estimation methods begin with this principle of using panel data to measure income changes across multiple frequencies to sort out the shocks.

[15]One interesting exception to the usual panel data approach is used by Blundell, Low and Preston (2013), who identify income shock variances in cross-section data using a combination of income and consumption data.

[16]There is a long-standing debate about whether estimated transitory variances are dominated by measurement error, which by construction will end up in the transitory shock terms. However, methodologically comparable estimates based on high-quality administrative data, such as in Sabelhaus and Song (2010), DeBacker et al. (2013) and Guvenen, Ozkan and Song (2014), are to a first approximation consistent with estimates from survey data, such as in Dynan, Elmendorf and Sichel (2007) and Gottschalk and Moffitt (2009).

The key questions about self-reported transitory income shocks in the SCF involve two high-level statistical properties of the gaps between actual and usual income: means and variances. The average gaps tend to be relatively small, though cyclical (Table 1). Variances of the self-reported gaps can be computed in a number of ways, but in order to have measures that are directly comparable to the residual-based estimates in the literature we compute the variance of the percentage gap using var[ln(actual income) – ln(normal income)]. Percentage gaps cannot be computed on zero or negative incomes, so we present two sets of estimates: the first has both actual and normal income restricted to be positive and the second has both restricted to be greater than $5,000.[17] For example, imposing the (modest) $5,000 threshold has a large impact on estimated variances; in 2013, the estimated variance falls from 11.4 per cent to 9.3 per cent.[18]

Given the statistical properties of the transitory component, the permanent or 'usual' component of income is much more stable than actual income, and thus a better classifier variable for looking within birth cohorts to compute changes over time.[19] The *level* of decomposition within birth cohort is determined by the precision with which the various statistics can be measured at given sample sizes. The analysis here is based on grouping families by usual incomes in the bottom half of the distribution, the 'next 45' per cent and the top 5 per cent. These populations are roughly equal in size across cohorts, because of the SCF oversampling strategy for high-wealth families.[20] For the purposes of our birth-year cohort-based figures, we assign a 'mid-point' age for each cohort in a given survey year to better compare across cohorts at any given point in the life cycle. This mid-point is always an age four years older than the minimum possible age for that cohort group in the given survey year – for example, in 2013, the mid-point age for the '1941–50' cohort was 67.

[17]In the 2010 SCF, only 0.5 per cent of families failed to meet the 'actual and normal income both greater than zero' condition and only 1.5 per cent failed to meet the $5,000 threshold.

[18]The same order of magnitude effect from imposing a lower bound on income has been observed in estimates of variances constructed using the residual method. See, for example, Sabelhaus and Song (2009 and 2010). Variance estimates in percentage terms are particularly sensitive to low initial values – an increase of income from $1,000 to $2,000 affects the estimated variance as much as a change from $100,000 to $200,000, though the two changes are obviously very different. Thus, one qualification for the assertion in the text that transitory variance estimates in the literature are roughly similar is that very small income values are effectively treated as zeros.

[19]The pitfalls of classifying families by actual income are well described in Bricker et al. (2014). In particular, the high-actual-income families suffering large negative transitory shocks in 2010 showed up with low actual incomes in the 2010 survey. They still had large wealth holdings when surveyed, of course, so they increased average holdings of assets such as corporate equities among 'low'-income families between 2007 and 2010.

[20]See the appendix to Bricker et al. (2014).

Fiscal Studies © 2016 Institute for Fiscal Studies

3. Education groups

In addition to looking within birth cohorts by usual income, we also apply the synthetic panel approach to groups based on the education of the head. Education is arguably an even more stable classifier than usual income, because the likelihood of changing education groups is very low after a certain age. However, the correlation between education and the economic characteristics in which we are interested (especially permanent income) is not as good, so the analysis does not speak as directly to the popular narratives about income, wealth and spending fluctuations across permanent income groups. As with usual income, we aggregate to three education groups (high school or less, some college, and college degree or higher) in order to achieve statistical precision on the within-cohort estimates. For most of the cohorts in the sample and at most ages, about 40 per cent of the population is in the first education group, 20 per cent in the second, and the remaining 40 per cent in the top education group. In general, the education-based analysis confirms the conclusions from the usual income groups in terms of the widespread nature of shocks and responses, though the decompositions of changes at young ages (where educational attainment does evolve in a predictable life-cycle way) are affected by movements across education groups.

V. Synthetic panel decomposition for income, wealth and spending on new cars

The overarching goal of this paper is to provide a set of facts that helps put various narratives about recent economic fluctuations in perspective. In an important sense, the estimates here are provided as statistical moments for calibrating the various theoretical constructs that have emerged to explain macroeconomic outcomes before, during and after the financial crisis. For example, a narrative based on the idea that credit liberalisation and credit tightening are the underlying *cause* of the boom and bust (as opposed to just correlated phenomena) should be confronted with the facts about whether or not the fluctuations are concentrated among the families for whom credit constraints were likely to be binding. Likewise, narratives based on inequality-driven fluctuations should be evaluated by looking at income, wealth and spending across the income distribution, in order to gauge *whose* behaviour changed the most.

1. Average incomes

The micro decomposition begins with average family incomes (Figure 6). This figure, and several to follow, provide the life-cycle perspective made possible

FIGURE 6

Mean income by birth-year cohort and usual income group

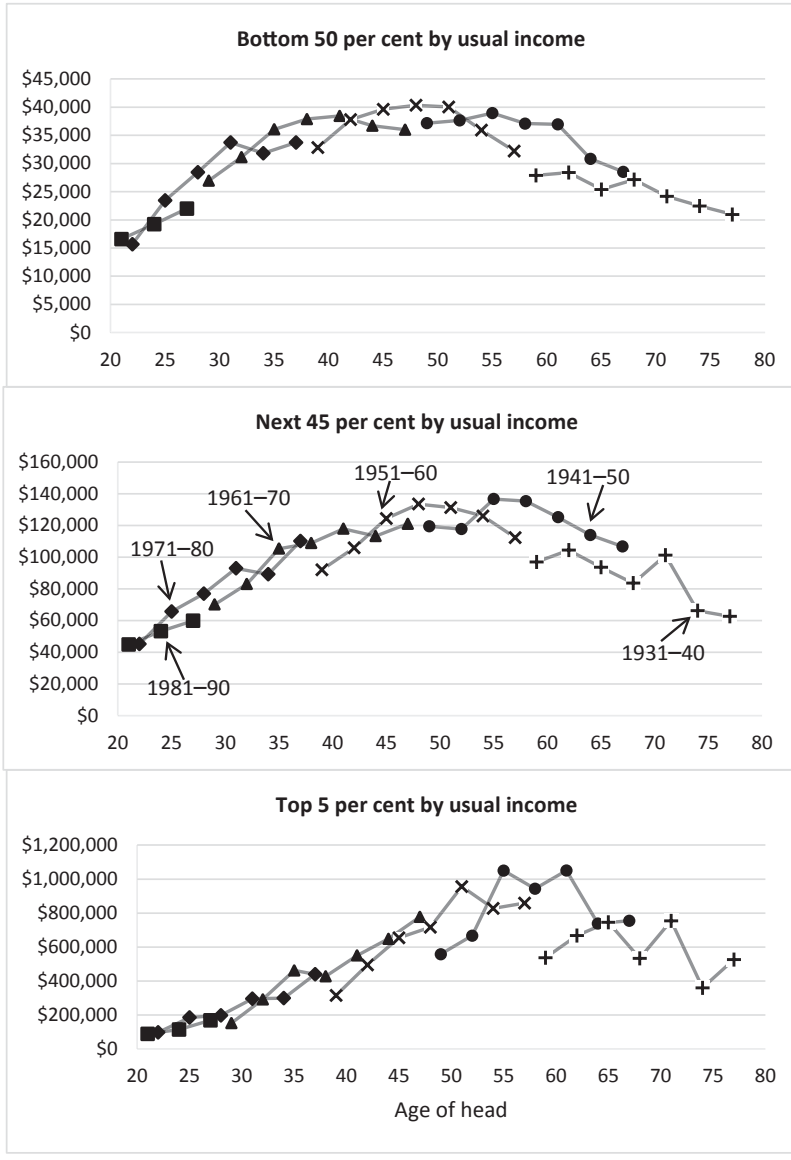

Note: Mean incomes are in 2013 dollars.
Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

by the synthetic panel approach. Each different line represents a birth cohort, for whom we observe outcomes at up to seven distinct points, three years apart, spanning up to 18 years of the life cycle. The average income values are in real terms, so any given cohort can be evaluated in terms of its own real growth trajectory, and relative to the cohorts ahead of or behind it (the vertical distance) where they overlap in given age ranges. The three panels of Figure 6 show average incomes across the three usual income groups: the bottom 50 per cent, the next 45 per cent and the top 5 per cent.

One important first observation when looking across the three life-cycle charts is the vertical scale, where the top of the average income range increases from $45,000 for the bottom 50, to $160,000 for the next 45 and $1.2 million for the top 5 per cent of families. Although the three usual income groups have very different levels of actual incomes, there are some common themes in the life-cycle trajectories. Most notably, the decline in average incomes in 2010 and generally continued low levels in 2013 (the last two observations for any given cohort line) relative to pre-recession trends are widespread. Almost every cohort / usual income group saw dramatic declines in average income during and after the Great Recession. Indeed, most groups saw their average incomes fall below the average incomes of the cohort 10 years ahead of them at the same age. In this sense, the focus of popular media on the plight of the young because of the Great Recession seems somewhat misplaced: the gap between realised incomes of (say) the 1951–60 birth cohort relative to the 1941–50 birth cohort in late middle age is much more dramatic.[21]

In addition to the broad similarities, there are a few notable differences across income groups in the life-cycle charts. The more rapid pre-crisis growth of actual incomes for higher usual income groups (especially the top 5 per cent) is widespread by age, and reflects the increasing income inequality during this period. Also, although there is only one post-recession observation, the young and top income groups have generally stabilised in terms of income levels, while the middle-aged and lower income groups have seen continued declines.

Decomposing the income distributions in 2007 and changes in income between 2007 and 2010 across birth cohort and usual income groups reinforces these visual impressions (Table 2). In 2007, the bottom half of families by usual income had 18 per cent of aggregate actual income, the next 45 per cent received 49 per cent and the top 5 per cent of families had a 33 per cent share. The shares of income changes between 2007 and 2010 (in this case, decreases) were more skewed towards the top, with the top 5 per cent accounting for 61 per cent of the income losses. Looking within the income changes panel, one sees income growth (indicated by negative signs, because total income is

[21] For a perspective on young adults' balance sheets in the aftermath of the financial crisis, see Dettling and Hsu (2014).

Fiscal Studies © 2016 Institute for Fiscal Studies

TABLE 2

Decomposing income changes by usual income groups and birth-year cohorts, 2007–10

Income shares, 2007

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	1%	3%	4%	5%	3%	1%	18%
Next 45	2%	7%	12%	14%	10%	5%	49%
Top 5	0%	3%	6%	11%	9%	4%	33%

Share of total income decrease, 2007–10

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	–6%	3%	0%	7%	8%	0%	12%
Next 45	–16%	–4%	4%	10%	18%	14%	27%
Top 5	–4%	–2%	–10%	21%	36%	21%	61%

Source: Survey of Consumer Finances.

falling) for the youngest age groups, and the most substantial income losses for cohorts in middle age in the top half of the usual income distribution.

Further support for the assertion that income shocks were widespread is found by looking within cohorts at education groups (Figure 7). The fact that education and income are correlated shows up in increasing ranges for the income scale as one moves from the 'high school or less' group to the 'college degree or higher' group, but the change in scales is much less dramatic than when looking directly at the usual income groups. The same basic observations come through clearly, however. Income shocks after 2007 were widespread, though relatively larger for middle age cohorts and more persistent for the low and middle education groups.

2. Average net worth

Shocks to household balance sheets during the financial crisis were even larger than shocks to income, and the synthetic panel analysis shows that these wealth shocks were widespread and persistent (Figure 8). The key driver of changes in household balance sheets over the period is of course asset prices, especially for owned housing, stocks, non-corporate business and other real-estate holdings. Differences in portfolio composition underlie the visual impression that wealth shocks were relatively larger for the bottom half of the usual income distribution, and that impression is confirmed by decomposing wealth shares and wealth change during the crisis (Table 3). In 2007 the bottom half by usual income owned 11 per cent of aggregate net worth, but between 2007 and 2010 the bottom half experienced 21 per cent of the aggregate loss.

The relative differential in wealth shocks across usual income groups during the recession years is explained by the dominance of owned housing in overall net worth for low- and modest-income families. House prices had only recently begun to rise when the 2013 SCF was conducted, while prices of financial assets had already largely rebounded. This helps to explain why average wealth in the bottom half continued to fall after 2010, while wealth generally stabilised for the top half (Figure 8). The same general observations – widespread wealth shocks followed by differential recoveries because of systematic differences in portfolio composition – hold when looking within birth cohorts by education (Figure 9). It is clear that higher-educated groups are well below pre-recession (life-cycle) trend in terms of average net worth, but the bottom half of families by usual income actually saw continued declines after 2010.

3. Spending on new cars

The observations on income and wealth shocks above are interesting in their own right, and the next set of questions involve whether those patterns can

FIGURE 7

Mean income by birth-year cohort and education group

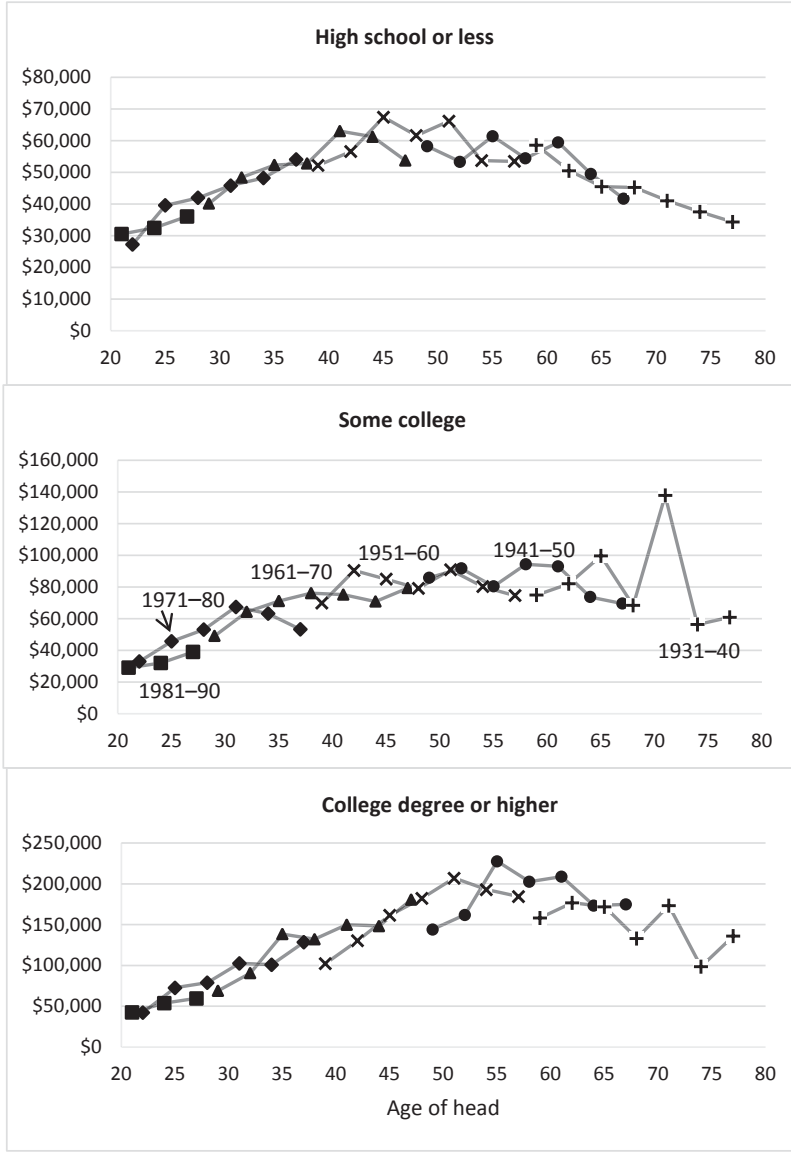

Note: Mean incomes are in 2013 dollars.
Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

FIGURE 8

Mean net worth by birth-year cohort and usual income group

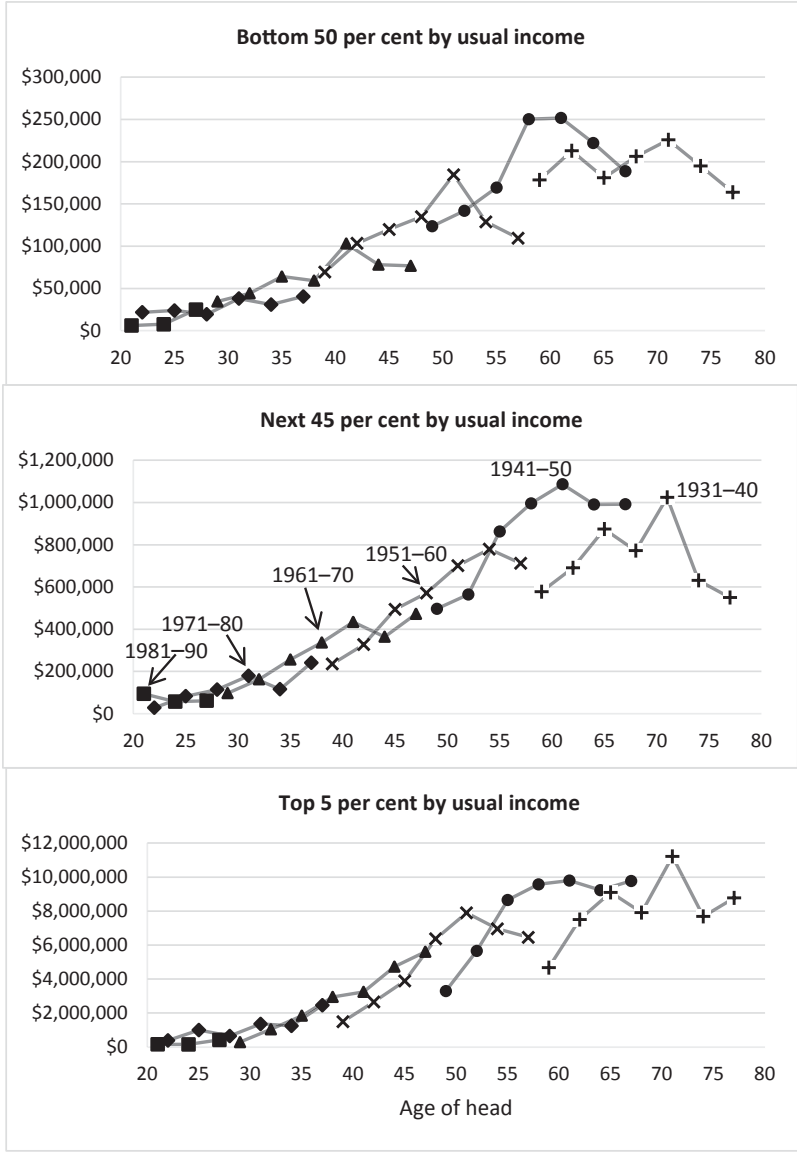

Note: Mean net worth values are in 2013 dollars.
Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

TABLE 3

Decomposing wealth changes by usual income groups and birth-year cohorts, 2007–10

Wealth shares, 2007

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	0%	1%	2%	4%	4%	2%	11%
Next 45	1%	2%	7%	12%	14%	7%	43%
Top 5	0%	2%	6%	15%	14%	9%	46%

Share of total wealth decrease, 2007–10

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	0%	1%	3%	10%	6%	1%	21%
Next 45	0%	6%	10%	–9%	19%	21%	47%
Top 5	–1%	1%	–22%	19%	14%	20%	32%

Source: Survey of Consumer Finances.

FIGURE 9

Mean net worth by birth-year cohort and education group

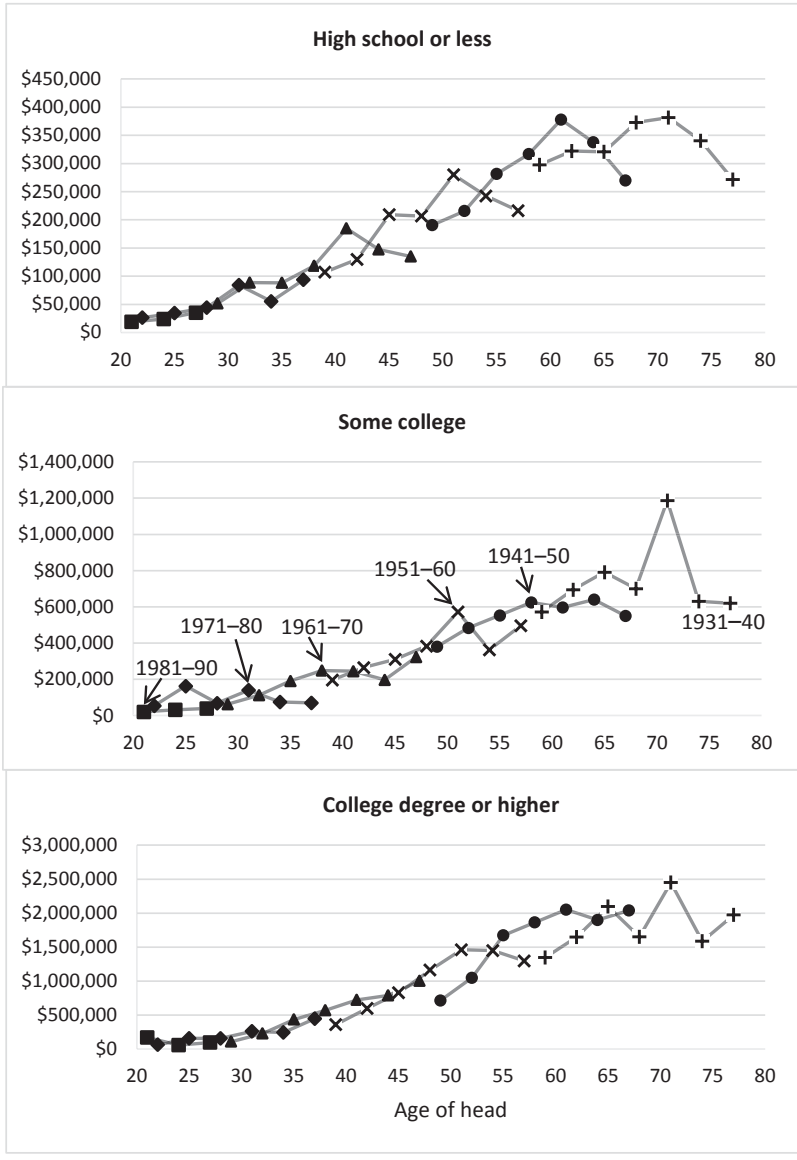

Note: Mean net worth values are in 2013 dollars.
Source: Survey of Consumer Finances.

improve our understanding of household spending during this period. The SCF does not collect a measure of overall spending, but the balance-sheet data on owned vehicles make it possible to track new car buying activity.[22] The specific measure we consider is the unconditional mean spent on new cars, because this measure includes both intensive and extensive decisions about car buying (Figure 10). The dramatic changes in car buying activity between 2007 and 2010, and continued reduced spending (as of 2013), show up clearly in the graphical synthetic panel analysis.

As with average income and wealth levels, it is important to keep in mind the vertical scales when looking at car spending across usual income groups. The top of the scale for (unconditional) average car spending is $3,500 for the bottom half of families by usual income, $10,000 for the next 45 per cent and $30,000 for the top 5 per cent of families. Thus, the top 5 per cent shows up as disproportionately accounting for 21 per cent of all car buying in 2007, the next 45 per cent accounts for 59 per cent and the bottom half only 19 per cent (Table 4, top panel).

The life-cycle trajectories for new car spending are much less hump-shaped than those for either income or wealth, and thus the time effects (visually) dominate the age effects (except perhaps for very young and very old age groups). This lack of life-cycle pattern makes the widespread nature of changes in car buying even more apparent, however, with virtually every birth cohort and income group exhibiting a pre-recession spike in car buying (unlike income and wealth, aggregate car buying was actually higher in 2004 than in 2007) followed by a substantial decline and persistently low (through 2013) level of spending. That visual 'widespread decline' impression shows up clearly in the decomposition, as (again) almost every group shows up as having reduced spending on new cars between 2007 and 2010 (Table 4, bottom panel).

Although 'widespread' is still the dominant impression one gets when looking across birth cohort and usual income groups in the 2007 to 2010 period, there is also some support for inequality-oriented narratives, because of the *relative* changes in behaviour at the bottom of the usual income distribution. The bottom half of families by usual income accounted for 19 per cent of new car spending in 2007, but they accounted for 36 per cent of the decline. The top 5 per cent of families accounted for 21 per cent of spending in 2007, but only 14 per cent of the decline between 2007 and 2010. One is left with a mixed impression about the inequality narrative, because the bottom half did exhibit a larger relative change, but the top half of families by usual income accounted for 64 per cent of the aggregate spending decline.

[22]The SCF can also be used to track new and used car buying activity, and the results are generally consistent. We focus on new car buying because it is the measure most closely associated with overall macroeconomic activity.

Fiscal Studies © 2016 Institute for Fiscal Studies

FIGURE 10

Mean spent on new cars by birth-year cohort and usual income group

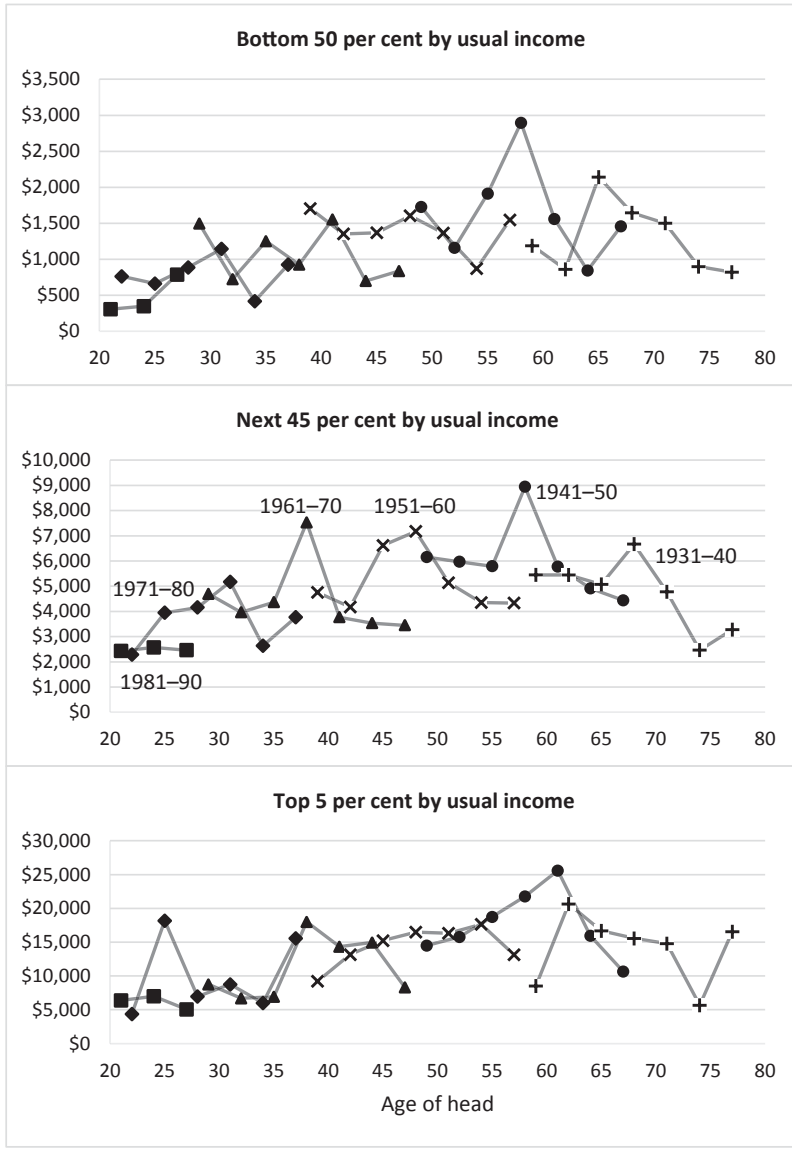

Note: New car spending is unconditional (includes zeros) and is shown in 2013 dollars.
Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

TABLE 4

Decomposing changes in new car spending by usual income groups and birth-year cohorts, 2007–10

Share of new car spending, 2007

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	0%	3%	5%	5%	4%	2%	19%
Next 45	3%	12%	10%	16%	13%	6%	59%
Top 5	1%	2%	4%	6%	6%	2%	21%

Share of new car spending decrease, 2007–10

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	–1%	9%	11%	7%	8%	3%	36%
Next 45	–7%	21%	2%	11%	11%	11%	49%
Top 5	–2%	3%	0%	–1%	11%	5%	14%

Source: Survey of Consumer Finances.

FIGURE 11

Mean spent on new cars by birth-year cohort and education group

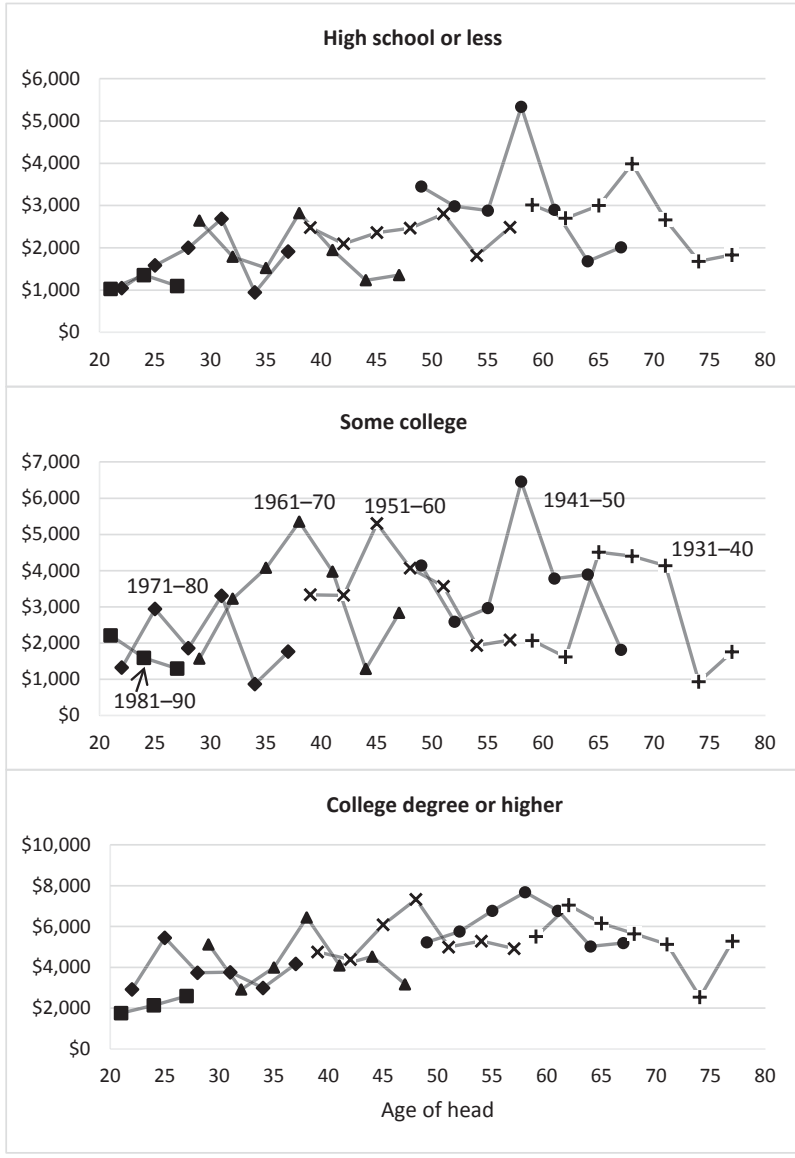

Note: New car spending is unconditional (includes zeros) and is shown in 2013 dollars.
Source: Survey of Consumer Finances.

The same basic patterns for new car spending show up when we look within birth cohorts by education (Figure 11). As with income and wealth, the vertical axes become much more compressed, because differences in car spending across education groups are much smaller than differences in spending across usual income groups. The story holds together quite well when education is used to create the synthetic panel, however, with (again) nearly every birth cohort and education combination exhibiting large drops in spending after 2007 (or even after 2004) and generally continued low spending on new cars through 2013.

VI. Housing and household debt

The inequality-related narratives that seek to explain recent macroeconomic fluctuations generally rely on owner-occupied housing as a key part of the story. The idea is that some change in the macroeconomic environment a couple of decades ago – often described generally as a 'loosening' of credit standards – expanded homeownership and/or housing-related debt for the subset of the population that was previously borrowing constrained. As house prices were booming, the cycle was self-fulfilling, and the lack of negative consequences for those increasing their housing-related borrowing led to even further increases in leverage, with credit happily provided by lenders who failed to fully consider the risks of extending credit to these borrowers. When the economic fundamentals changed – especially house price dynamics – those same families then had difficulty meeting their new debt obligations, defaults ensued, and the impact on lender balance sheets triggered the financial crisis associated with the Great Recession. Subsequently, those families have seen slow income growth and been denied credit, worsening their already impaired balance-sheet and cash-flow situation.

The synthetic panel approach is well suited to address these features of the inequality narrative. As with the income and wealth shock / spending analysis in the previous section, there are some observations about homeownership, housing turnover and debt that support the narrative, but the preponderance of the evidence suggests much more widespread changes in behaviour. The supporting evidence comes from the fact that low- and modest-income families borrowed more *relative to their ex post income growth* than families at the top of the usual income distribution, but the counter-argument (as with income, wealth and spending) focuses on the decomposition of macroeconomic aggregates. In particular, the surge in housing turnover and debt in the boom period and collapse in housing transactions and borrowing during the bust was widespread, highlighted by the fact that the share of aggregate debt outstanding across permanent income groups changed very little.

Fiscal Studies © 2016 Institute for Fiscal Studies

1. Homeownership

The first housing-related measure to consider using the synthetic panel approach is homeownership itself (Figure 12). The top of each scale is 100 per cent, and the systematic life-cycle patterns are clear, as the trajectories for all three usual income groups are upward sloping with age, and homeowning is nearly universal in the top 5 per cent and is above 90 per cent in the 'next 45' income group for families with heads over age 30. Only the bottom half of the usual income distribution exhibits homeownership rates below 50 per cent through middle age, with peaks (for pre-1950 cohorts) approaching 70 per cent later in life.

The questions raised by the inequality-related narrative involve comparing across cohorts within income groups, however. Here, there is some evidence of accelerated movements into homeownership during the housing boom. For example, the homeownership rate of the 1971–80 birth cohort, 'next 45' usual income group was about 75 per cent when they were in their early 30s in 2007. The same income group in the previous cohort (birth years 1961–70) had a homeownership rate almost 10 percentage points lower at the same age 10 years earlier. However, that sort of accelerated homeownership attainment is not observed for the bottom half of the income distribution, whose life-cycle trajectories were very stable through the boom years.

The more striking results in the homeownership charts are for the housing bust and subsequent periods. There is clear evidence that overall homeownership within birth-cohort groups in the bottom half of the distribution was thrown off the existing trajectories. After 2007, every lower-income cohort born after 1950 failed to achieve the levels of homeownership reached by the cohort ahead of them at the same age 10 years earlier. In some groups, such as the 1951–60 cohort, the data suggest absolute declines in homeownership rates after 2007.

2. Recent home buying activity

Although the cross-cohort comparisons of homeownership rates do not suggest dramatic changes in behaviour by particular groups during the housing boom, our measure of housing turnover does indicate striking fluctuations (Figure 13). The concept being measured here is the fraction of all families in a given birth cohort / usual income group that reported having transacted the owned home in which they are currently living within the past three years. Although this is not a measure of housing investment per se, it is highly correlated with new residential construction, and is a preferable measure in most ways because we are interested in how families were changing in terms of their own housing investment and assumption of debt obligations.

The denominator of the housing turnover measure is all families, so transactions will naturally be higher at certain ages and for income groups

FIGURE 12

Homeownership rate by birth-year cohort and usual income group

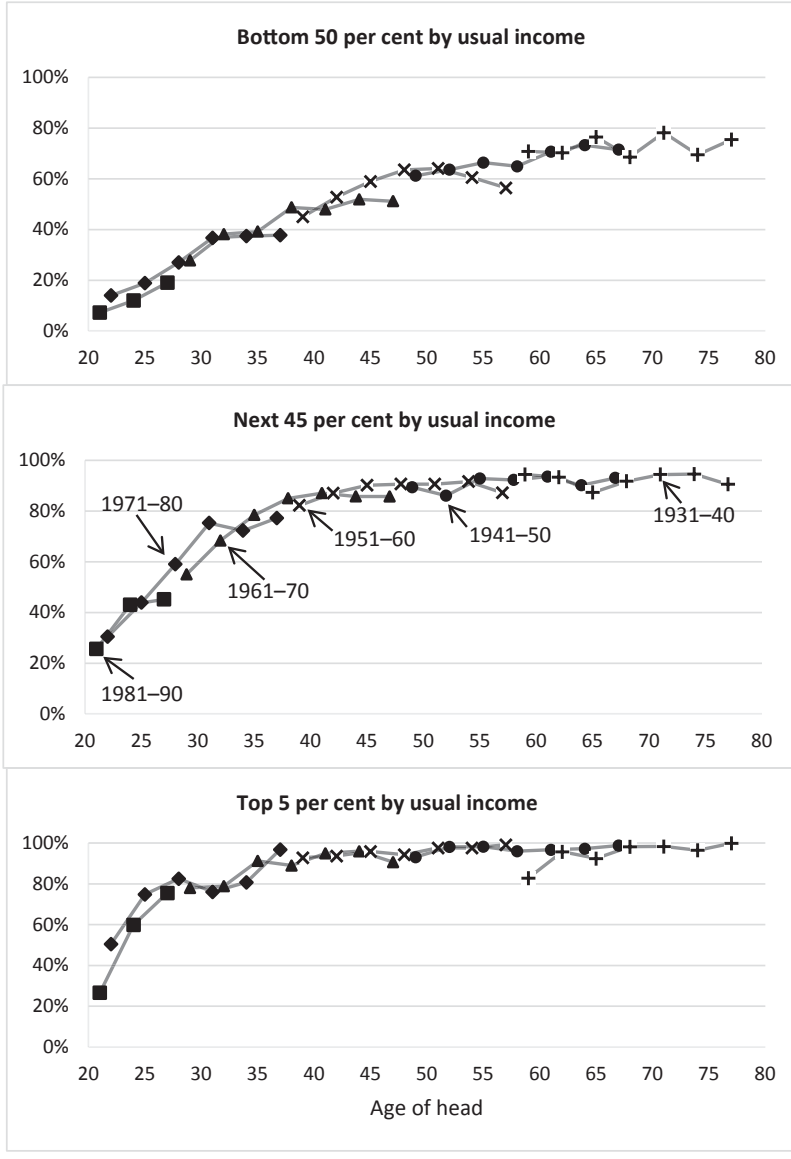

Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

FIGURE 13

Percentage buying new primary residence in the last three years by birth-year cohort and usual income group

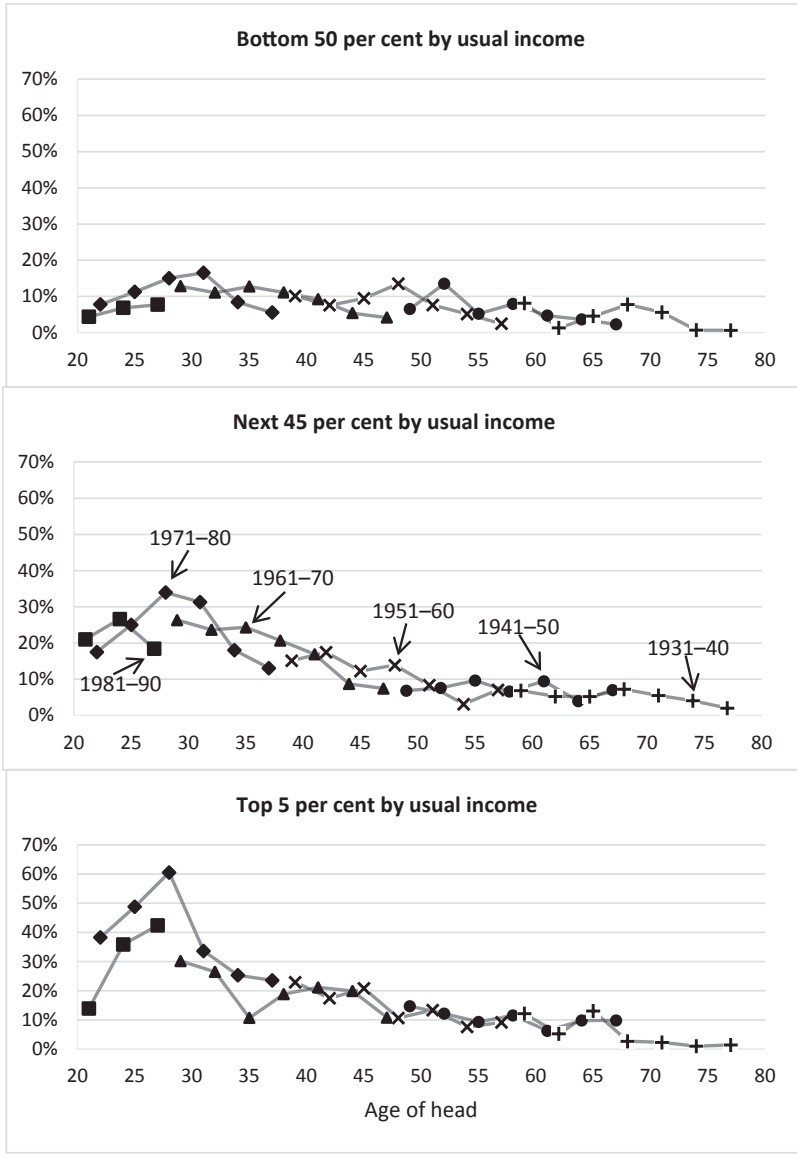

Note: Based on owner-reported date of acquiring currently-owned residence.
Source: Survey of Consumer Finances.

who are entering into housing. Indeed, the peak of the transaction measure is not quite 20 per cent for the bottom half of the usual income distribution in their early thirties, while for the top 5 per cent of the 1971–80 birth cohort, some 60 per cent of families reported a transaction in their late twenties. As with homeownership itself, however, the signal we are looking for is from the across-cohort differences, and here there is ample evidence that during the housing boom – especially in the top half of the usual income distribution – families were transitioning housing at a much greater pace than the cohorts ahead of them at the same age. This is especially true within the 'next 45' usual income group, the same group that post-crisis saw the biggest declines in housing transactions, even when viewed from the life-cycle perspective.

3. Household debt

In many ways, the role of housing in the inequality narrative is to help us understand the dramatic increase in household debt during the boom, most of which was housing-related. We consider the growth of debt in the life-cycle framework using two approaches, focusing on trajectories of debt to income (Figure 14) and the decomposition of aggregate debt and changes in debt during the boom and bust (Table 5). As with the other synthetic panel measures, there is some evidence in support of the inequality narrative, as debt to income ratios rose less during the boom for the highest-income families. However, that result is driven largely by the fact that incomes were growing more rapidly at the top, which may or may not have been anticipated when debt levels were rising. The overall distribution of outstanding debt across usual income groups was actually little changed during the boom, and may have even shifted toward higher-income families during the bust.

The life-cycle view of the synthetic panel trajectories shows fairly widespread increases in debt relative to income across age and income groups during the boom (Figure 14). Again, this is reflected by the fact that younger cohort trajectories lay above the preceding group, with the gaps widening between the 2001 and 2007 surveys. This is true for virtually all age groups in the bottom 95 per cent of the usual income distribution, and all but the middle-aged and older groups in the top 5 per cent. However, the differential growth is not concentrated in the bottom half of the usual income distribution, as the changes for the next 45 per cent and for the younger age groups in the top 5 per cent were all noticeably above the previous cohort's.

The widespread growth in debt relative to income shows up in the overall decomposition of debt and debt changes for the boom and bust periods (Table 5). The bottom half of the usual income distribution was responsible for 21 per cent of the debt outstanding in 2001, while the next 45 per cent accounted for 60 per cent of debt and the 'top 5' group for the remaining 19 per cent. The growth of debt across usual income groups in the boom

Fiscal Studies © 2016 Institute for Fiscal Studies

FIGURE 14

Mean debt to mean actual income by birth-year cohort and usual income group

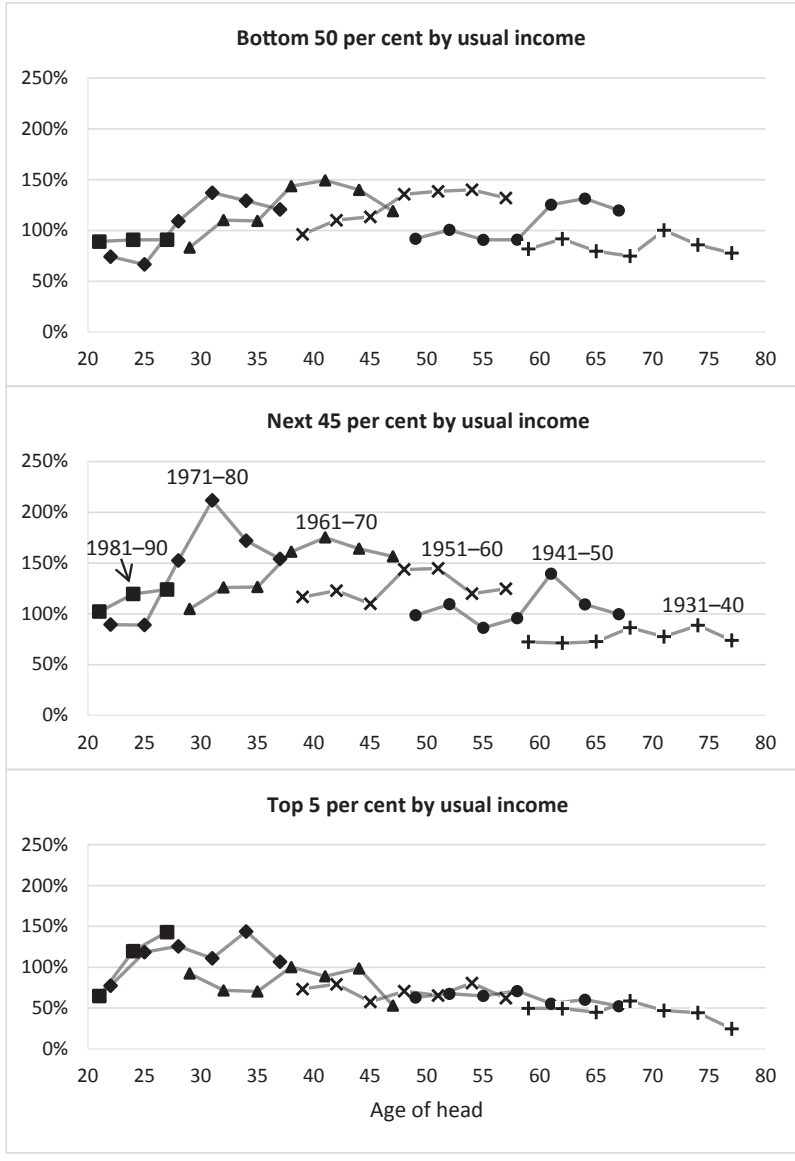

Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

TABLE 5

Decomposing changes in debt by usual income groups and birth-year cohorts, 2001–07 and 2007–13

Share of debt, 2001

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	0%	2%	6%	8%	4%	2%	21%
Next 45	0%	6%	16%	20%	12%	5%	60%
Top 5	0%	2%	4%	6%	5%	2%	19%

Share of debt growth, 2001–07

	Birth cohort						
Usual income	*1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	1%	7%	5%	3%	2%	0%	19%
Next 45	4%	24%	17%	12%	10%	–1%	67%
Top 5	1%	3%	5%	5%	1%	0%	15%

TABLE 5

(Continued)

Share of debt, 2007

Usual income	Birth cohort *1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	1%	4%	5%	6%	3%	1%	20%
Next 45	2%	13%	17%	17%	11%	2%	62%
Top 5	0%	3%	4%	6%	4%	1%	18%

Share of debt decline, 2007–13

Usual income	Birth cohort *1981–90*	*1971–80*	*1961–70*	*1951–60*	*1941–50*	*1931–40*	*All*
Bottom 50	–7%	1%	11%	14%	7%	3%	28%
Next 45	–28%	1%	9%	38%	41%	6%	66%
Top 5	–13%	–14%	6%	13%	8%	6%	6%

Source: Survey of Consumer Finances.

Fiscal Studies © 2016 Institute for Fiscal Studies

was quite similar, however, leaving 2007 debt shares largely unchanged from 2001. The decline in debt after 2007 was also widespread, though families in the bottom half accounted for a disproportionate share of debt reduction (28 per cent), presumably because some debt was discharged in mortgage defaults or access to new credit was impaired. The top 5 per cent income group actually accounted for the least amount of relative deleveraging, at 6 per cent of the total decline in debt after 2007.

VII. Conclusions

Life-cycle consumption theory provides us with a number of different ways to explain the dramatic fluctuations in household spending that have occurred in the past two decades. Indeed, in some ways, the problem with life-cycle theory is that it provides *too many* different ways to explain the same spending patterns, and choosing between the potential causal factors and behavioural mechanisms is the great challenge. Distinguishing between the various causal relationships is of great importance to public policy, because our understanding of what went wrong in the past is our guide to improving policy in the future.

The analysis here is focused on providing the sort of facts that will hopefully improve our collective ability to disentangle the various fundamentals driving fluctuations in household spending over the most recent business cycle. The facts do not validate narratives focused on changes in income, wealth and spending for the specific population subgroups for whom rising inequality and access to credit markets are most problematic, though these groups did (as in any turbulent macroeconomic times) experience amplified outcomes. The data suggest that shocks to income and wealth, and corresponding changes in spending behaviour, were widespread across the population in the period leading up to, during and after the Great Recession. It would have been impossible to have had such a spending boom, a severe downturn and such a slow recovery, were it not for the fact that the top half of the permanent income distribution also experienced shocks and changed their behaviour.

The facts presented here suggest a reinterpretation or at least a substantial modification of certain views about what transpired in the period leading up to, during and beyond the Great Recession. The boom and bust in household borrowing and spending was not simply driven by policies expanding access to homeownership and credit for previously under-served groups. The situation is better described as one in which there was widespread belief that the economic fundamentals driving house prices, borrowing and increased spending were actually sound. When beliefs about house prices and other economic fundamentals were proven wrong, and widespread behavioural changes ensued, the families who became unable to meet their financial obligations experienced the brunt of the collateral damage.

Fiscal Studies © 2016 Institute for Fiscal Studies

Supporting information

Additional supporting information may be found in the online version of this paper on the publisher's website:

- Data Appendix

References

Attanasio, O. and Borella, M. (2014), 'Movements in individual consumption: a time-series analysis of grouped data', *International Economic Review*, vol. 55, pp. 959–91.

Berger, D. and Vavra, J. (2015), 'Consumption dynamics during recessions', *Econometrica*, vol. 83, pp. 101–54.

Bertrand, M. and Morse, A. (2013), 'Trickle-down consumption', National Bureau of Economic Research (NBER), Working Paper no. 18883.

Blundell, R., Low, H. and Preston, I. (2013), 'Decomposing changes in income risk using consumption data', *Quantitative Economics*, vol. 4, pp. 1–37.

Bricker, J., Dettling, L. J., Henriques, A., Hsu, J. W., Moore, K. B., Sabelhaus, J., Thompson, J. and Windle, R. A. (2014), 'Changes in U.S. family finances from 2010 to 2013: evidence from the Survey of Consumer Finances', *Federal Reserve Bulletin*, vol. 100, no. 4, pp. 1–40.

—, Henriques, A., Krimmel, J. and Sabelhaus, J. (2015), 'Measuring income and wealth at the top using administrative and survey data', Federal Reserve Board, Finance and Economics Discussion Series (FEDS) Working Paper no. 2015-30.

—, Ramcharan, R. and Krimmel, J. (2014), 'Signaling status: the impact of relative income on household consumption and financial decisions', Federal Reserve Board, Finance and Economics Discussion Series (FEDS) Working Paper no. 2014-76.

Carroll, C. D. (1992), 'The buffer-stock theory of saving: some macroeconomic evidence', *Brookings Papers on Economic Activity*, no. 2, pp. 61–156.

Chetty, R. and Szeidl, A. (2007), 'Consumption commitments and risk preferences', *Quarterly Journal of Economics*, vol. 122, pp. 831–77.

Christelis, D., Georgarakos, D. and Jappelli, T. (2015), 'Wealth shocks, unemployment shocks and consumption in the wake of the Great Recession', *Journal of Monetary Economics*, vol. 72, pp. 21–41.

Crossley, T. F. and Low, H. W. (2014), 'Job loss, credit constraints, and consumption growth', *Review of Economics and Statistics*, vol. 96, pp. 876–84.

Cynamon, B. Z. and Fazzari, S. M. (2014), 'Inequality, the Great Recession, and slow recovery', Federal Reserve Bank of St. Louis, Working Paper.

Deaton, A. (1985), 'Panel data from time series of cross sections', *Journal of Econometrics*, vol. 30, pp. 109–26.

DeBacker, J., Heim, B., Panousi, V., Ramnath, S. and Vidangos, I. (2013), 'Rising inequality: transitory or permanent? New evidence from a panel of U.S. tax returns', *Brookings Papers on Economic Activity*, Spring, pp. 67–142.

De Nardi, M., French, E. and Benson, D. (2012), 'Consumption and the Great Recession', *Economic Perspectives*, vol. 36, pp. 1–16.

Dettling, L. J., Devlin-Foltz, S. J., Krimmel, J., Pack, S. J. and Thompson, J. P. (2015), 'Comparing micro and macro sources for household accounts in the United States: evidence from the Survey of Consumer Finances', Federal Reserve Board, Finance and Economics Discussion Series (FEDS) Working Paper no. 2015-86.

— and Hsu, J. W. (2014), 'The state of young adults' balance sheets: evidence from the Survey of Consumer Finances', *Federal Reserve Bank of St. Louis Review*, vol. 96, pp. 305–30.

Dynan, K. E., Elmendorf, D. W. and Sichel, D. E. (2007), 'The evolution of household income volatility', Federal Reserve Board, Finance and Economics Discussion Series (FEDS) Working Paper no. 2007-61.
Gottschalk, P. and Moffitt, R. (2009), 'The rising instability of U.S. earnings', *Journal of Economic Perspectives*, vol. 23, no. 4, pp. 3–24.
Guvenen, F., Ozkan, S. and Song, J. (2014), 'The nature of countercyclical income risk', *Journal of Political Economy*, vol. 122, pp. 621–60.
Jappelli, T. and Pistaferri, L. (2014), 'Fiscal policy and MPC heterogeneity', *American Economic Journal: Macroeconomics*, vol. 6, no. 4, pp. 107–36.
Justiniano, A., Primiceri, G. E. and Tambalotti, A. (2015), 'Household leveraging and deleveraging', *Review of Economic Dynamics*, vol. 18, pp. 3–20.
Kaplan, G. and Violante, G. L. (2014), 'A model of the consumption response to fiscal stimulus payments', *Econometrica*, vol. 82, pp. 1199–239.
—, — and Weidner, J. (2014), 'The wealthy hand-to-mouth', *Brookings Papers on Economic Activity*, Spring, pp. 77–138.
Kumhof, M., Rancière, R. and Winant, P. (2015), 'Inequality, leverage, and crises', *American Economic Review*, vol. 105, pp. 1217–45.
Mian, A. and Sufi, A. (2014a), 'House price gains and U.S. household spending from 2002 to 2006', National Bureau of Economic Research (NBER), Working Paper no. 20152.
— and — (2014b), 'What explains the 2007–2009 drop in employment?', *Econometrica*, vol. 82, pp. 2197–223.
Moffitt, R. (1993), 'Identification and estimation of dynamic models with a time series of repeated cross-sections', *Journal of Econometrics*, vol. 59, pp. 99–123.
Petev, I., Pistaferri, L. and Eksten, I. S. (2012), 'Consumption and the Great Recession: an analysis of trends, perceptions, and distributional effects', in D. Grusky, B. Western and C. Wimer (eds), *The Great Recession*, New York: Russell Sage Foundation.
Piketty, T. and Saez, E. (2003), 'Income inequality in the United States, 1913–1998', *Quarterly Journal of Economics*, vol. 118, pp. 1–39.
— and — (2015), 'The World Top Incomes Database', http://topincomes.parisschoolofeconomics.eu/.
Sabelhaus, J. and Ackerman, S. (2012), 'The effect of self-reported transitory income shocks on household spending', Federal Reserve Board, Finance and Economics Discussion Series (FEDS) Working Paper no. 2012-64.
— and Song, J. (2009), 'Earnings volatility across groups and time', *National Tax Journal*, vol. 62, pp. 347–64.
— and — (2010), 'The great moderation in micro labor earnings', *Journal of Monetary Economics*, vol. 57, pp. 391–403.
Saez, E. and Zucman, G. (2014), 'Wealth inequality in the United States since 1913: evidence from capitalized income tax data', National Bureau of Economic Research (NBER), Working Paper no. 20625.

Fiscal Studies © 2016 Institute for Fiscal Studies