FISCAL STUDIES, vol. 41, no. 1, pp. 7–38 (2020) 0143-5671

Consumption and Wage Inequality in the US: The Dynamics of the Last Three Decades*

NEIL AMIN-SMITH† and ORAZIO P. ATTANASIO‡

†*Institute for Fiscal Studies*
(neil.amin.smith@gmail.com)
‡*Yale University; Institute for Fiscal Studies; FAIR @ NHH; National Bureau of Economic Research*
(orazio.attanasio@yale.edu)

Abstract

In this paper, we look at the evolution of consumption and wage inequality from 1980 to 2016 in the US. We use data from the Consumer Expenditure Survey (CEX) and the Current Population Survey (CPS) to look at differences in consumption and wages across groups in the population defined by educational attainment of the household head and year-of-birth cohort. We show that the results obtained by Attanasio and Davis (1996) for non-durable consumption still hold in more recent decades. In addition to non-durable consumption and services, we look at inequality measured in terms of expenditure on and stock of vehicles. The advantages of looking at these measures are that information on cars is typically measured more accurately than other components of expenditure and consumers are more likely to react by adjusting their stock of vehicles on the basis of long-term expectations about their economic prospects.

I. Introduction

Discussion of inequality in the US and elsewhere has tended to focus on rising disparities in income – it is generally recognised that income inequality, at

* Submitted June 2019.

The authors would like to thank Peter Levell for much help with the CEX data and, in particular, the most recent waves of the CEX. Both Amin-Smith's and Attanasio's time was partly funded by the Economic and Social Research Council through the Centre for Microeconomic Analysis of Public Policy at IFS (ES/M010147/1).

Keywords: consumption inequality, income shocks, insurance.

JEL CODE: D12, D31, E21, I31, D63, J31, J24.

least as measured by official statistics, has increased substantially over the past half-century. The importance attached to this trend arises partly from an interpretation of inequality in income as being a useful proxy for inequality in economic well-being more broadly.

There are, however, a number of reasons why changes in income inequality might not accurately capture changes in the inequality of long-run resources that households or individuals have access to.[1] At the simplest level, official income statistics do not take account of taxes and transfers, but this is possible to rectify. Harder to disentangle is the fact that income measures may reflect transitory changes in income rather than changes in the distribution of people's lifetime income, which may be more relevant for economic well-being. For example, those experiencing temporary income volatility may be able to offset such fluctuations by borrowing or saving or might receive countervailing transfers from family members.[2] In principle, even permanent shocks could, *ex ante*, be insured. In addition, measures of income inequality may miss the effects of financial or other wealth, or debt burdens, on well-being – such effects could well be independent of changes in income.

As recognised by Blundell and Preston (1998) among others, consumption may be a better guide to expected lifetime resources than income, and it is for this reason that it may be more informative to study consumption patterns to learn about changes in economic well-being than it is to study income patterns. Furthermore, the simultaneous study of income and consumption fluctuations can also be informative about the ability of individual households to insure certain types of shocks and, therefore, the types of formal and informal markets individuals have access to. This approach is pursued in Blundell, Pistaferri and Preston (2008) and, more recently, extended to consider female labour supply by Blundell, Pistaferri and Saporta-Eksten (2016). This framework has been extended by Blundell, Pistaferri and Saporta-Eksten (2018) to consider children and by Attanasio, Meghir and Mommaerts (2019) to consider extended families. Arellano, Blundell and Bonhomme (2017) have moved away from (log) linear models of income and considered more complex models of earnings in conjunction with consumption.

Motivated by these insights, a number of papers have also looked specifically at consumption inequality in the US, mostly using the Consumer Expenditure Survey (CEX), to examine whether trends in consumption inequality have mirrored those in wage or earnings inequality. The findings of earlier work were mixed, although some studies suggested that the

[1] Or, more generally, changes in inequality in well-being.

[2] If a certain market structure is assumed, then it becomes possible to decompose transitory and permanent shocks to income, which is key in literature that aims to test consumption smoothing and the permanent income hypothesis – see, for example, Blundell, Pistaferri and Preston (2008).

increase in consumption inequality towards the end of the 20^{th} century and in the early years of the 21^{st} was relatively modest by comparison with the increase in income inequality.[3] More recently, mounting evidence that there is considerable non-classical measurement error in the CEX has resulted in a number of papers that employ various methods to address this issue. These papers have tended to find that consumption inequality in the US since the 1980s has tracked increases in income inequality relatively closely.[4]

In this paper, we contribute to the literature that has looked at consumption (and income) inequality in three ways. First, we document trends in the distribution of both wages and income on one hand and consumption on the other across birth cohorts and education groups between 1980 and 2016, contributing to knowledge on how patterns of inequality in economic well-being have evolved over that period.

Second, we frame the relationship between relative wage movements and relative household consumption across groups in terms of tests of risk sharing models. In particular, we frame the study using perfect insurance as a benchmark. Risk sharing may happen through a variety of explicit and implicit insurance mechanisms that insulate the distribution of consumption from shocks to the distribution of income. The literature that this paper thus contributes to aims to test whether observed consumption corresponds to the predictions of such models – that growth in household consumption should be cross-sectionally uncorrelated with idiosyncratic variables that are also uncorrelated with consumers' preferences. Unsurprisingly, previous work has tended to reject full consumption insurance[5] but the underlying theory can still provide a useful benchmark against which to examine the extent to which income shocks are reflected in shocks to consumption. Our analysis updates that of Attanasio and Davis (1996), who study how relative hourly wage movements across birth cohorts and education groups affected the distribution of household consumption and economic welfare, by extending the data to 2016. The extended time frame for the sample allows us to add 25 years to the original 10-year sample and also incorporates the Great Recession. As in Attanasio and Davis, we test the full insurance hypothesis over a range of specifications, each emphasising co-movements of different frequencies. We also quantify the extent to which our results are driven by education and by cohort separately.

Third, we look at the evolution of inequality in the stock of and expenditure on cars. Cars are particularly attractive for our analysis for several reasons.

[3] See, for example, Krueger and Perri (2006).

[4] Attanasio, Battistin and Ichimura, 2007; Meyer and Sullivan, 2010; Attanasio, Hurst and Pistaferri, 2014; Aguiar and Bils, 2015.

[5] For example, Cochrane (1991), Mace (1991), Altonji, Hayashi and Kotlikoff (1992), Townsend (1994) and Dynarski and Gruber (1997).

First, the quality of the micro data we use is very high. There are virtually no issues in replicating aggregate data sources by aggregating the individual data on car expenditure (by contrast, there is substantial evidence that it is difficult to match figures from National Income and Product Accounts by aggregating CEX data[6]). Second, cars are very salient to individual choices: they are large and adjusting the value of their stock can be costly. Decisions on car stocks and car purchases, therefore, are likely to reflect individual expectations of income for several years into the future. Finally, if the utility from cars and the utility from non-durable consumption are non-separable, tests of risk sharing based on variation of the marginal utility of consumption should take such non-separability into account. Thus we derive an augmented test of risk sharing that accounts for this non-separability using data on households' stocks of and expenditures on cars.

Section II describes the data that we make use of in this paper. Section III documents relative consumption, income and wage changes across cohort–education groups between 1980 and 2016 as well as describing the recent evolution of inequality in both car ownership and expenditures on car purchases.

Section IV begins with a perfect insurance model and derives implications in terms of changes in marginal utility. Based on these results, we detail the test of perfect insurance that we are able to carry out by using synthetic panel data to examine the impact of systematic shifts in the hourly wage structure on the distribution of household consumption. We also discuss specification issues that might arise from preference non-separabilities, and set out how we address these by including durables (in particular, cars) as a non-separable argument of the utility function.

In Section V, we carry out the tests of perfect insurance described in the previous section at different frequencies – which serves to provide some indicative evidence relating to consumption smoothing and the permanent income hypothesis. In Section VI, we augment our tests of perfect insurance to account for possible non-separability of preferences with regards to durables.

II. The data

Following Attanasio and Davis (1996), our empirical analysis makes use of consumption data from the Consumer Expenditure Survey (CEX) and income data from the Current Population Survey (CPS), both of which are large survey data sets collecting information on individuals and households in the United States.

The CEX collects data on expenditure, income and demographic characteristics of consumer units – which are defined as a group of individuals

[6]See McCarthy et al. (2002).

living in the same household who are related or share at least two of three major expenditure categories. It surveys approximately 5,000 households each year. The CEX is in fact made up of two distinct surveys – the diary survey and the interview survey. We make use in this paper of the interview survey, in which respondents are sampled every three months for a total of four quarters. We construct a measure of consumption that is equal to household expenditure on non-durable goods and services, which excludes expenditure on durables, health, education and housing. As in Attanasio and Davis (1996), we exclude non-urban households, those residing in student housing, those with a male head older than 59 or younger than 23, and those with incomplete income responses.[7] The CEX samples each household four times, so that each household reports up to 12 monthly observations. When a household does not have the complete 12 observations *in a given calendar year*,[8] we compute annual consumption measures weighting each household in proportion to the number of monthly observations that fall into that calendar year, taking into account seasonal effects.

The CPS collects data on household income, demographic characteristics and the labour market outcomes of individual members. This information is collected for roughly 40,000–60,000 households per year – hence the advantage of using income data from the CPS rather than the CEX to construct our relative wage measures. In addition to the substantially larger sample size, CPS income and earnings data are superior to the corresponding CEX data.[9] We construct a measure of hourly earnings from the CPS, computed as annual earnings divided by the product of number of weeks worked and usual hours per week. We apply the same age restrictions as for the CEX, and also exclude persons who were students or in the military for at least part of the year and those who earned less than 75 per cent of the minimum wage.

In both sets of data, we form synthetic panel groups by combining cohort groups with educational attainment groups, and from a sample restricted to households with a male head. Birth cohorts are defined in terms of five-year bands, and we consider four educational attainment categories: fewer than 12 years of schooling, exactly 12 years, more than 12 but fewer than 16 years, and 16 or more years of schooling.

[7]Younger heads and those in student housing are excluded to minimise migration to higher education attainment categories as a cohort ages. Excluding old heads minimises the impact of retirement and retirement choices on our sample. We exclude non-urban households because the CEX did not sample them in 1982 or 1983. While this introduces a potential bias, we follow common practice here, both because it is not possible to correct the sample and because we do not think it would make a big difference, given the size of the excluded population.

[8]Either because the household does not respond to all four interviews or because the monthly observations span two different years.

[9]Cutler and Katz, 1991.

III. Changes in relative wages, incomes and consumption

1. Inequality in income and relative wages

In Figures 1 and 2, we report how inequality in (equivalised) total post-tax family income[10] changed between 1980 and 2016. Figure 1 plots the standard deviation of log equivalised total family income, which shows a substantial increase across the period, with three periods of particularly sharp increase in the early 1980s, the first half of the 1990s and since 2008. As shown on the graph, increases in the early 1980s and increases between 2008 and 2010 coincided with periods of recession in the US economy. After an increase in the early 1990s, there was little sustained change in the standard deviation until the Great Recession – the standard deviation in 2008 was even slightly lower than it was in 1996. Figure 2 shows further how the distribution of income changed over this period – in particular, it shows the ratios between various percentiles of the distribution of family income in each year. The ratio between the 90th percentile and the 10th percentile exhibits very similar trends to those shown in Figure 1 – sharp increases in the early 1980s, the early 1990s and since 2008. The 50:10 ratio also exhibits a similar increase in the early 1980s,

FIGURE 1

Standard deviation of log equivalised total family income

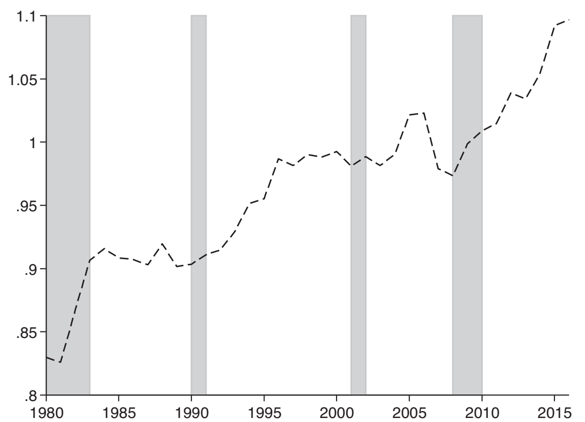

Note: Grey bars denote when the US economy was in recession (rounded to the nearest year), according to the National Bureau of Economic Research (NBER) Business Cycle Reference Dates. Total family income is defined as described in Section II and in footnote 10, and equivalised according to Betson (2004)'s widely-supported update to the National Academy of Sciences (NAS) recommended equivalence scale.
Source: Authors' calculations using the CPS, various years.

[10]For this we use the variable *ftotval* from the CPS. The questionnaire text defines this as 'money from jobs, net income from business, farm or rent, pensions, dividends, interest, social security payment and any other money income received by members of this family who are 15 years or over'.

FIGURE 2

Ratios of percentiles of log equivalised total family income

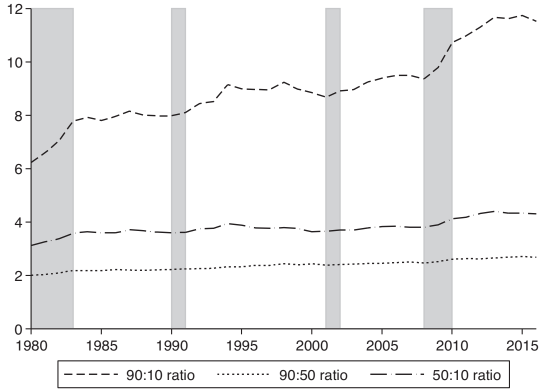

Note: As for Figure 1.
Source: As for Figure 1.

FIGURE 3

Real hourly wage movements for men by birth cohort and education group

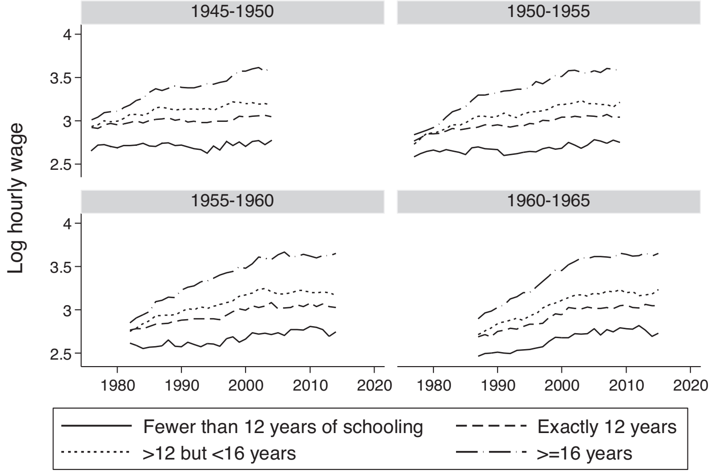

Note: Each graph describes real wage movements for a different birth cohort.
Source: Authors' calculations using the CPS, various years.

and remained mostly flat from 1995 until 2008, after which it began to increase once more. The ratio between the 90th percentile and the 50th percentile instead shows a steadier increase across the whole period.

We also document the evolution of inequality in relative real wages. In particular, Figure 3 shows the evolution of real hourly wages for the four cohorts that are in our sample for the longest period of time (since we exclude from our sample those aged under 23 or over 59). What can be clearly seen is that the gap between those with 16 or more years of education and those with fewer widened over time in each of the cohorts shown.

Table A1 in the online appendix shows real hourly wage movements for men between 1980 and 2015 by birth cohort and education group across all cohorts. The vast majority of cohort–education groups are shown to have experienced rising average real wages over the period we are considering. However, across all cohorts presented, differences by educational attainment are shown to have widened markedly during this period. This result is consistent with what is reported by Attanasio and Davis (1996) for the period between 1980 and 1990. To take the 1955–60 birth cohort as an example, real wages among college-educated men grew by 41 log points relative to those with fewer than 12 years of education between 1985 and 2010. The real wage gap between those with some post-secondary education and those with fewer than 12 years of education can also be shown to have increased across all cohorts.

2. Inequality in non-durable consumption

In Figure 4, we report the standard deviation of log equivalised household consumption between 1980 and 2015. We observe a sharp increase in inequality in the early 1980s, which mirrors that shown with regards to income in Figure 1. After that early period, however, there is a much slower gradual drift upwards. Since 2000, the figure shows that the standard deviation has remained relatively

FIGURE 4

Standard deviation of log equivalised household non-durable consumption

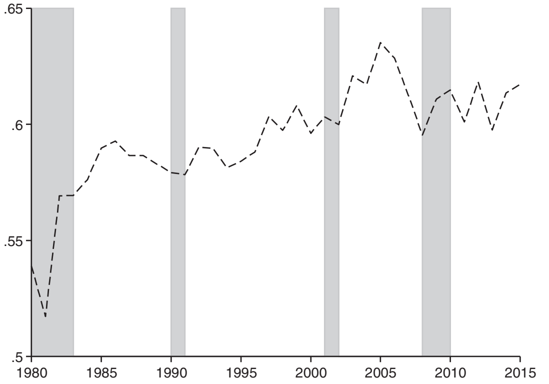

Note: Measure of consumption constructed as described in Section II. Equivalisation as for Figure 1. Grey bars denote periods when the US economy was in recession, as for Figure 1.
Source: Authors' calculations using the CEX, various years.

FIGURE 5

Ratios of percentiles of log equivalised household non-durable consumption

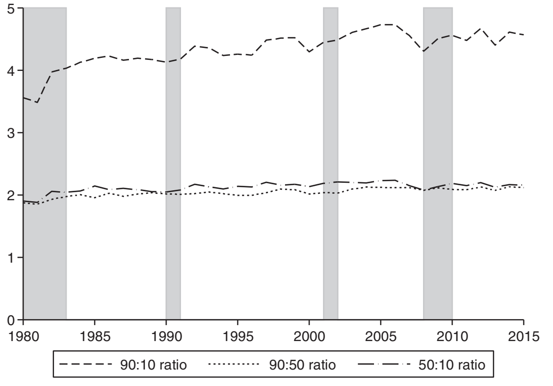

Note: As for Figure 4.
Source: As for Figure 4.

flat, and certainly does not appear to reflect the spikes in income inequality shown in the early 1990s and since 2008. Notably, both of the prolonged periods of recession coincide with increases in the standard deviation of consumption – the early 1980s and 2008–10.

We also show, in Figure 5, how household consumption at different points of the distribution evolved between 1980 and 2015. There appears to be considerably less dispersion in the consumption distribution than in the income distribution. For example, the 90:10 ratio in 2015 was nearly 2.5 times higher for household income than for household consumption.

Furthermore, across the whole period, the 90:10 ratio in consumption increases by around 30 per cent compared with a near doubling of the corresponding ratio for the income distribution. There does appear to have been a more consistent increase in the ratio between consumption at the 90th and 10th percentiles than there was in the standard deviation of consumption, although this again increases most sharply at the very beginning of the period.

Although we appear to show a substantially larger increase in the 90:10 ratio for consumption than some other work using the same data – for example, Meyer and Sullivan (2017) – the differences are not as stark as they first look. In fact, a large part of the increases in the standard deviation and the 90:10 ratio between 1980 and 2015 that are shown in Figures 4 and 5 took place in the early 1980s, a period that Meyer and Sullivan do not cover.[11] Furthermore, broadly

[11] This is because between 1981 and 1984 the survey only includes respondents from urban areas. Meyer and Sullivan thus leave out these years, whilst, as mentioned in footnote 7, we restrict our sample in all years to urban households in order to construct continuous but consistent measures.

FIGURE 6

Real non-durable consumption for households by birth cohort and education group of household head

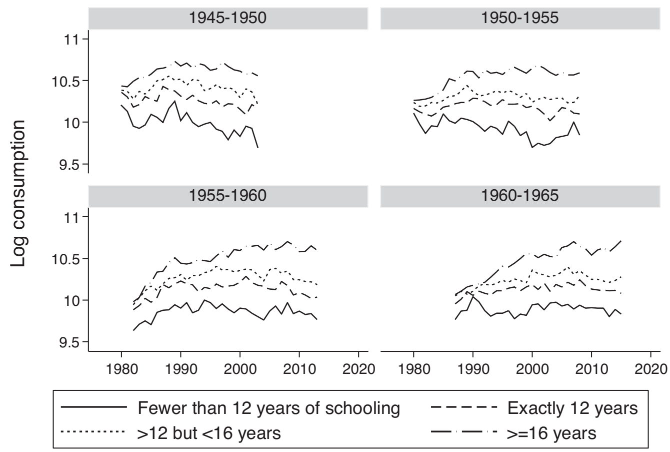

Note: Each graph describes movements in real non-durable consumption for a different birth cohort.
Source: Authors' calculations using the CEX, various years.

speaking, our findings reflect their work in that both the standard deviation and the 90:10 ratio of household consumption have been relatively flat compared with the evolution of income inequality over the same period.

The 90:50 ratio also shows a steady, but much slighter, increase over the whole period, whilst the ratio between the 50th and 10th percentiles increases in the early 1980s but shows little overall change thereafter – it was at a similar level in 2015 to that in 1985 (although there are some mild fluctuations shown in between, including a shallow but sustained decrease during the second half of the 1980s). This can be compared with a 50:10 ratio in the income distribution that increases considerably across the whole period.

Figure 6 shows how the evolution of real non-durable consumption has differed between education groups in the cohorts in our sample for the longest. As with real wages, it is clear that the gap between those with 16 or more years of education and those with fewer grew across the period. However, it can also be seen that consumption for all four education groups grew by considerably less across the period than real wages did.

Table A2 in the online appendix describes the evolution of real non-durable consumption between 1980 and 2015 for all cohort groups. It shows that growth in real consumption has differed dramatically between education groups. Taking the 1955–60 birth cohort as an example once more, real non-durable household consumption for those with college education rose by 21 log points relative to those with fewer than 12 years of education between 1985 and

2010. Among some of the older cohorts, there are in fact substantial falls in the household consumption of those with fewer than 12 years of education. This on its own suggests that the consumption insurance hypothesis is unlikely to hold.

Of course, there are a number of issues in using CEX data on total consumption to track inequality in consumption.[12] Whilst earlier work often found that consumption inequality did not track changes in inequality in income over the period as a whole (as shown in Figures 4 and 5), more recently researchers employing various strategies to overcome issues of measurement error in the CEX have refuted this – finding much more substantial increases in consumption inequality over the period in question than those shown here. Aguiar and Bils (2015), for example, make use of the relative expenditures of high- and low-income households on luxuries versus necessities to correct for measurement error and find that increases in consumption inequality mirror those in income inequality to a much greater extent than implied by reported total expenditure. Attanasio, Hurst and Pistaferri (2014) employ a number of methods to overcome measurement error, including focusing on consumption categories where measurement error has been found to be less of an issue. They find that consumption inequality between 1980 and 2010 increased by nearly the same amount as income inequality.

3. Decomposition of inequality in income and consumption

In order to explore further the evolution of inequality in income and consumption documented above, we construct a Theil index for each year and exploit its decomposability. Theil indices are a particular case of generalised entropy measures, and allow one to decompose inequality into the part that is due to between-group inequality and the part that is due to within-group inequality. In particular, the Theil index is defined as

$$(1) \qquad \text{Theil} = \frac{1}{N} \sum_{i=1}^{N} \frac{x_i}{\mu} \ln\left(\frac{x_i}{\mu}\right),$$

where x_i is the income or consumption of household i and μ is the mean of x. This can be decomposed such that

$$(2) \qquad \text{Theil} = \sum_{c=1}^{m} s_c T_c + \sum_{c=1}^{m} s_c \ln\left(\frac{\overline{x_c}}{\mu}\right),$$

where s_c is the income share of group c, $\overline{x_c}$ is the average income or consumption of households in group c, and T_c is the Theil index for group c. The first term

[12] A discussion of some of these issues can be found in Attanasio and Pistaferri (2016).

FIGURE 7

Theil index of inequality in household income

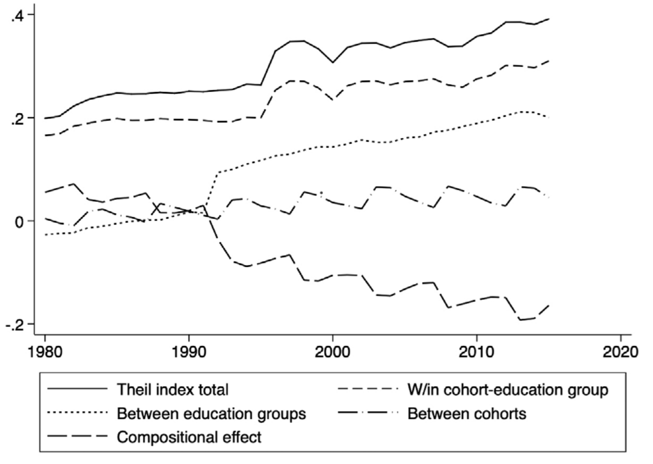

Note: As for Figure 1.
Source: As for Figure 1.

can then be seen to represent within-group inequality and the latter term to represent between-group inequality.

Building on this, we are able to decompose the Theil index for each year into four components. Simple decomposition allows us, for each year, to separate out variation between households into inequality between cohort–education groups and inequality within cohort–education groups. Having done this, we are able to further decompose the first of these terms to separately identify inequality between cohorts (this will capture both age and cohort effects), inequality between education groups, and a component reflecting the composition of our sample in each year with regards to cohort and education.

Figure 7 shows the evolution of these components for household income between 1980 and 2015. The decomposition shows that the increase in the Theil index is mirrored by a substantial increase in the element due to inequality between education groups, as well as an increase in inequality within cohort–education groups. By contrast, the part of inequality relating to the composition of the population in each year with regards to cohort and education is shown to reduce over time. This may reflect the increase over this period in the proportion of people with higher levels of education.

Figure 8 shows the evolution of the Theil index for household consumption between 1980 and 2015, decomposed in the same way. Whilst the Theil index for household consumption is shown to have remained relatively flat across this period, it can again be seen that there is a rise in inequality between

FIGURE 8

Theil index of inequality in household consumption

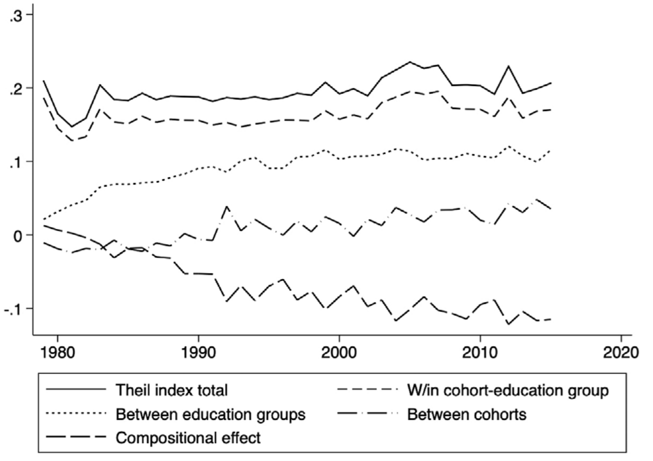

Note: As for Figure 4.
Source: As for Figure 4.

education groups. Once again, the component due to the composition of the population declines over time, whilst the within-group element closely tracks overall inequality – remaining relatively flat. This last observation suggests that there does seem to be some within-group consumption insurance, since Figure 7 showed that within-group income inequality rose across the period. Between education groups, however, the trend in inequality in consumption seems to reflect that in income – suggesting that there is little consumption insurance between education groups.

4. Inequality in stock of and expenditure on cars

Figure 9 shows how the mean value of the stock of vehicles held by households has changed over time (in 2016 prices). It shows an increase of nearly 20 per cent between 1984 and 2004 – from \$10,000 to a peak of around \$12,000. Between 2007 and 2010, this increase was reversed, with stocks falling to just above \$9,500, almost as low as they had been at any point over the period shown. Notably, around the time of recessions in the early 1990s and 2008–10, the mean vehicle stock fell – in both cases, this decline began before the economy entered recession. As would be expected, these falls were not overly dramatic, but their existence is unsurprising given the cyclical nature of car purchases.[13] The lack of a decline in stocks during the 2001 recession reflects

[13] See, for example, Greenspan and Cohen (1996).

FIGURE 9

Mean equivalised stock of vehicles held by households

Note: Equivalisation as for Figure 1.
Source: Authors' calculations using the CEX, various years.

what others have noted – that purchases of durables held up uniquely well through that particular recessionary period.[14]

Over the same period, inequality in car ownership, as measured by the standard deviation, is shown in Figure 10 to have increased substantially. The figure shows that it increased sharply around the mid 1980s and again from around 1997. A steep fall between 2004 and 2008, however, brought its level by the time of the Great Recession back to where it had been in the mid 1990s.[15]

We are also able to document expenditure on cars during this period. Figure 11 shows the mean value of car purchases per household between 1980 and 2015. It shows that, whilst mean car expenditures per household doubled between 1980 and 2002, there was a sustained decline between 2002 and 2009 in overall (net) car expenditure that more than reversed this, with some limited recovery since then.

Figures 12 and 13 attempt to decompose this pattern – showing, respectively, the percentage of households purchasing a car for each year and the average net value of car purchases conditional on positive net expenditure on cars in that year. Bar-Ilan and Blinder (1992) note that fluctuations in car purchases tend to be driven by changes in the number of people buying cars rather than by the average expenditure per purchase. However, whilst Figure 12 does show that the percentage of households buying at least one car fell in the

[14]For example, Leamer (2007).

[15]We also plotted the standard deviation of households' residualised car stocks, regressing on region, state, age of household head, number of children and number of adults in the household. The resulting graph was very similar to Figure 10.

FIGURE 10

Standard deviation of households' equivalised car stocks

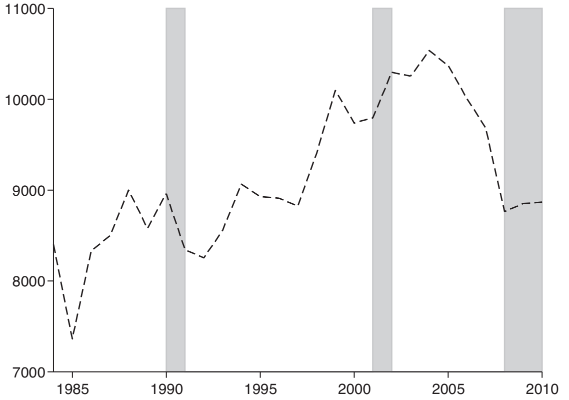

Note: As for Figure 9.
Source: As for Figure 9.

FIGURE 11

Mean (equivalised) car purchase expenditure per household

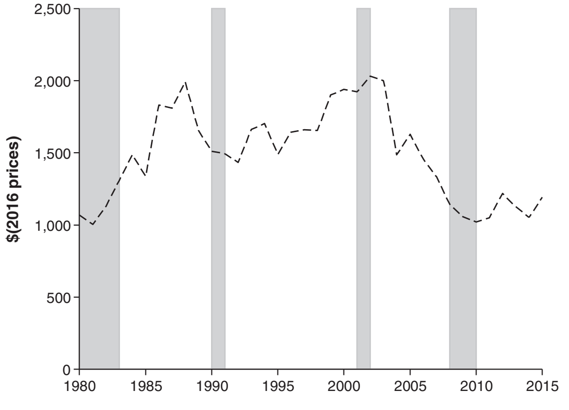

Note: Expenditure is net expenditure – expenditure on car purchases minus any proceeds from trading in cars. Equivalisation as for Figure 1.
Source: Authors' calculations using the CEX, various years.

two years leading up to the Great Recession, that percentage had been falling since the late 1980s. Furthermore, Figure 13 shows that households' average expenditure conditional on purchasing was falling between 2007 and 2010. Whilst an investigation into this pattern is beyond the scope of this paper, explanations put forward have included the collapse of credit (both generally and in affecting the financing of purchases by car manufacturers) and the size and persistence of the shocks experienced during the Great Recession.

FIGURE 12

Percentage of households buying at least one car

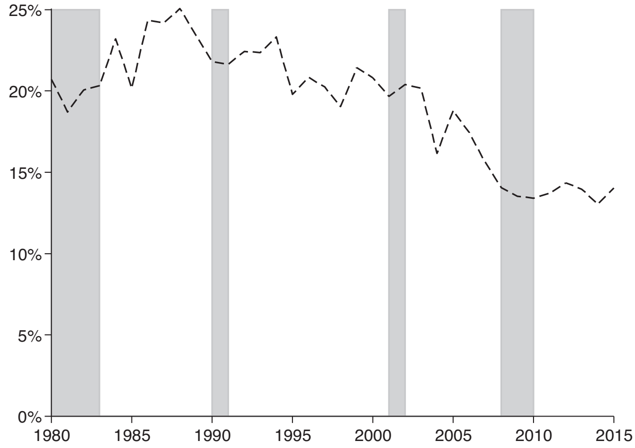

Source: Authors' calculations using the CEX, various years.

FIGURE 13

Mean net value of household car purchases conditional on purchasing a car

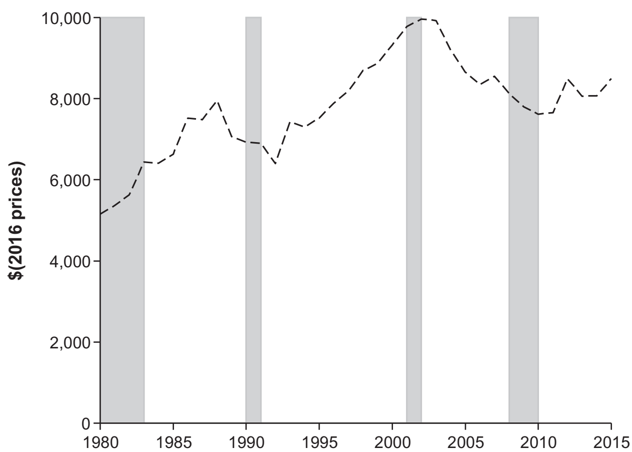

Note: As for Figure 11.
Source: As for Figure 11.

Finally, Figure 14 shows inequality in (net) car expenditures among those with positive net expenditure on cars in each year. Although the pattern is relatively flat across the period, the coefficient of variation is mostly over 100 per cent, implying significant variance in car expenditures between households.

FIGURE 14

Coefficient of variation of car expenditures (conditional on purchasing in that year)

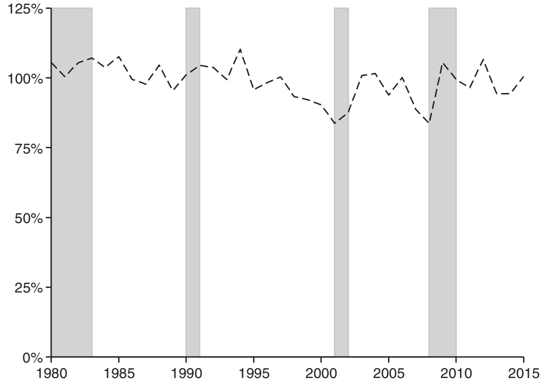

Note: As for Figure 11.
Source: As for Figure 11.

IV. Are relative changes in wages reflected in consumption changes? A conceptual framework

Changes in the between-group elements of inequality demonstrated by the Theil decompositions in the previous section provide obvious motivation for investigating between-group insurance more systematically. In this section, therefore, we derive the tests that enable us to do this.

As mentioned in the introduction, tests of consumption insurance consider an important implication of perfect insurance – namely, that in such an environment, growth in household consumption should be cross-sectionally uncorrelated with idiosyncratic variables that are also uncorrelated with consumers' preferences. Such an implication can be derived from the first-order condition of a planner's problem that considers the optimal allocation of resources across individuals and over time, given a fixed set of Pareto weights given to the individual households. In particular, it is assumed that, given N individuals and an infinite horizon, at time t the planner maximises the following expression:

$$(3) \qquad \max \sum_{i=1}^{N} \pi_i \sum_{k=0}^{\infty} \beta^k E_t[U(c_{i,t+k}, z_{i,t+k}, \omega_{i,t+k})]$$

$$\text{s.t. } S_{t+k+1} = (1 + R_{t+k})S_{t+k} + \sum_{i=1}^{N} y_{i,t+k} - \sum_{i=1}^{N} c_{i,t+k},$$

where π_i is the Pareto weight for household i, β is the discount factor, and $c_{i,t}$ and $y_{i,t}$ are consumption and income for household i at time t. S is consumers' wealth and R the rate of return it generates. $z_{i,t}$ is a vector of observable variables that affect utility, while $\omega_{i,t}$ is a vector of unobservable variables; either of these two sets of variables can be a choice variable (although, for notational simplicity, we are not including them in the budget constraint). A first-order condition for the planner problem in equation 3 is

$$(4) \qquad \pi_i \beta^k \frac{\partial U}{\partial c_{i,t+k}} = \mu_{t+k} \; \forall \, k \geq 0,$$

where μ_{t+k} denotes the Lagrange multiplier associated with the aggregate feasibility constraint at time $t+k$. The marginal utility $\frac{\partial U}{\partial c_{i,t+k}}$ may depend on $z_{i,t+k}$ and/or $\omega_{i,t+k}$.

As in this version of the model there is full information and insurance contracts are fully enforceable, it should be noted that equation 4 does not hold in expectation but for every possible history up to time $t+k$ and every possible realisation of the state variables that govern the evolution of the economy. Empirically, that implies that any 'residual term' one would include when bringing equation 4 (or its log version) to the data would reflect either measurement error or unobserved (to the econometrician) taste shifters such as $\omega_{i,t+k}$. We also notice that the right-hand side of the equation does not depend on the individual index i. That is the key implication of full risk sharing, as the only relevant constraint is that of aggregate resources.

1. Level and difference specifications

If we consider equation 4 for $k = 0$ and $k = 1$, and take the ratio of the two expressions, we get

$$(5) \qquad \beta \frac{\frac{\partial U}{\partial c_{i,t+1}}}{\frac{\partial U}{\partial c_{i,t}}} = \frac{\mu_{t+1}}{\mu_t} \equiv \nu_{t+1}.$$

Assuming power utility, so that $U\left(c_{i,t}, z_{i,t}, \omega_{i,t}\right) = \frac{(c_{i,t})^{1-\gamma}}{1-\gamma} \exp(\delta' z_{i,t} + \omega_{i,t})$, and taking the log of the resulting expression, equation 4 becomes

$$(6) \qquad \ln\left(\pi_i\right) + k \ln\left(\beta\right) - \gamma \ln\left(c_{i,t+k}\right) + \delta' z_{i,t+k} + \omega_{i,t+k} = \ln\left(\mu_{t+k}\right),$$

which, when being considered for $k = 0$ and $k = 1$, subtracting one from the other, yields

$$(7) \qquad \ln\left(\beta\right) - \gamma\left[\ln\left(c_{i,t+1}\right) - \ln\left(c_{i,t}\right)\right] + \delta' \Delta z_{i,t+1} + \Delta\omega_{i,t+1} = \Delta\ln\left(\mu_{t+1}\right).$$

When considering equation 7, we notice that it does not need to be written in first differences (as it is), as one would get a similar expression even when considering non-adjacent time periods. In this case, equation 7 becomes

$$(8) \qquad k \ln(\beta) - \gamma \Delta_k \ln(c_{i,t+k}) + \delta' \Delta_k z_{i,t+k} + \Delta_k \omega_{i,t+k} = \Delta_k \ln(\mu_{t+k}).$$

For empirical purposes, it is convenient to rewrite equations 6 and 8 as follows:

$$(9) \qquad \ln(c_{i,t+k}) = \varphi_i + \tilde{\delta}' z_{i,t+k} + \tilde{\mu}_t + \tilde{\omega}_{i,t+k}$$

$$(10) \qquad \Delta_k \ln(c_{i,t+k}) = \kappa + \tilde{\delta}' \Delta_k z_{i,t+k} + \Delta_k \tilde{\omega}_{i,t+k} - \Delta_k \ln(\mu_{t+k}).$$

Equation 4 and its corresponding equation 9 state that log consumption for individual i depends on a fixed effect φ_i, which reflects the discount factor and the individual Pareto weight in the social planner problem, on observable and unobservable taste shifters $z_{i,t+k}$ and $\tilde{\omega}_{i,t+k}$ and on a time fixed effect $\tilde{\mu}_t$ that does not depend on the index i.

By considering different time periods, equation 5 eliminates the fixed effect φ_i and states that the planner allocates resources so that discount-weighted growth in marginal utility is equal across individuals. Equations 8 and 10, which correspond to equation 5 for a specific utility function, imply that any variable that is cross-sectionally uncorrelated with preference variation and measurement error in consumption growth is also uncorrelated with the cross-sectional distribution of consumption growth. Therefore, they imply that if a variable is cross-sectionally uncorrelated with preference variation and measurement error, then consumption insurance implies that that variable has no explanatory power for the cross-sectional distribution of consumption growth.

These equations reflect the essence of the perfect insurance hypothesis, under which changes in log marginal utility are constant in the cross section as individuals diversify idiosyncratic risk. The equations and the model are silent on the specific mechanisms through which risk is diversified and focus on the cross-section dynamics of consumption allocation. Tests of consumption insurance are thus based around regressions of this form.

In particular, the literature has often considered these specifications:

$$(11) \qquad \ln(c_{i,t+k}) = \varphi_i + \tilde{\delta}' z_{i,t+k} + \tilde{\mu}_t + \lambda \ln(y_{i,t+k}) + \tilde{\omega}_{i,t+k}$$

$$(12) \qquad \Delta_k \ln(c_{i,t+k}) = k + \tilde{\delta}' \Delta_k z_{i,t+k} + \tilde{\lambda} \Delta_k \ln(y_{i,t+k}) + \Delta_k \tilde{\omega}_{i,t+k},$$

where perfect insurance implies that $\lambda = 0$ and $\tilde{\lambda} = 0$. In equations 11 and 12, the variable $y_{i,t+k}$ typically represents individual income or some sorts of shocks to individual resources which are assumed not to affect the marginal utility of consumption. Such shocks to resources, once one controls for aggregate time effects and (in equation 11, which is specified in levels) individual fixed effects, should not affect the level of or changes in the marginal utility of consumption, which, conditional on the taste shifters $z_{i,t+k}$, is proxied by log consumption.

We note that analysing different frequencies (by changing the k in Δ_k) allows us to construct tests of different power under different hypotheses. Suppose the log of individual income $y_{i,t+k}$, expressed in deviation from group levels, is made of two components – temporary and permanent – as is often assumed in the literature:

$$\ln(y_{i,t+k}) = v^p_{i,t+k} + u_{i,t+k} \tag{13}$$

$$v^p_{i,t+k} = v^p_{i,t+k-1} + \xi_{i,t+k}, \tag{14}$$

where $\xi_{i,t+k}$ and $u_{i,t+k}$ are i.i.d.[16] over time and across individuals and represent the permanent and temporary shocks to individual resources respectively. Such an income process implies

$$\Delta_k \ln(y_{i,t+s}) = \sum_{j=0}^{k-1} v^p_{i,t+s-j} + \Delta_k u_{i,t+s}. \tag{15}$$

When analysing first differences (or higher frequencies), the test focuses on short-run shocks and gives more weight to temporary shocks than when the value of k is relatively high. In this latter case, the measure of idiosyncratic shocks to resources gives more weight to permanent shocks and relatively little weight to temporary shocks. The same applies to the levels equations, if one expresses the level of income as the accumulation of all past permanent shocks. If temporary shocks are easier to smooth (for instance, via dis-saving) than the permanent ones, the insurance tests that rely more on variation induced by permanent shocks might be more powerful at detecting deviations from first-best and perfect insurance.

2. Non-separabilities

The presence of taste shifters such as $z_{i,t+k}$ makes the approach very flexible and potentially powerful. Such variables could reflect demographic effects (such as changing needs induced by changes in family composition), but could also reflect non-separability of non-durable consumption and other determinants of utility, such as leisure or durables. The theory delivers implications for

[16]Independent and identically distributed.

the marginal utility of consumption and, therefore, if non-separabilities of consumption with other determinants of utility are important, they should be reflected in the tests of perfect insurance.[17] At the same time, however, the presence of such variables makes the identification of deviations from perfect risk sharing particularly difficult, especially if the parameter vector $\tilde{\delta}$ is not known. A good example of such difficulty is the case where one considers real wages as a potential variable that should appear in equation 11 or 12. If one allows the marginal utility of consumption to depend on leisure, one ends up with an equation that relates (log) consumption, (log) leisure and (log) wages, which can then be interpreted as the first-order condition of a labour supply problem that relates the marginal rate of substitution between consumption and leisure to real wages.

In what follows, we consider two potentially important non-separabilities: one in which we let the marginal utility of consumption depend on female labour supply and the other where we let it depend on the stock of cars held by the household. We will consider the identification issues just discussed when we interpret the results obtained with these extensions.

Preferences that are non-separable between consumption and leisure could lead to false rejections of the consumption insurance hypothesis, in particular by giving rise to a correlation between measured consumption growth and relative wage movements. In order to control for non-separability between consumption and women's leisure, we augment our benchmark specification with the group mean of log female leisure hours. This is aimed at controlling for shifts in the marginal utility of household consumption. As mentioned above, the presence of non-separabilities can make testing full insurance difficult. In our case, while we allow female leisure to affect the marginal utility of consumption, we identify shocks to resources using male wages. The implicit assumption is that male leisure is separable from consumption and, therefore, under perfect risk sharing, changes in male wages should not be reflected in equations such as 11 or 12.

Another possible specification issue arises from the fact that the measure of consumption we use to test the perfect insurance hypothesis is consumption of *non-durable* goods and services. This means that if preferences are non-separable between consumption of non-durables and consumption of durables, the marginal utility of non-durable consumption, and therefore equations 11 and 12, should include among the *taste shifters* the stock of durables. Looked at another way, the log of non-durable consumption is no longer proportional to

[17] As we discuss below, the lack of longitudinal data forces us to consider pseudo panels, which we obtain by aggregating equations such as 11 and 12 or 11′ and 12′ across the individuals belonging to a predefined group. This approach forces us to work with expressions that are linear in parameters, as will be apparent below. Therefore, when considering non-separabilities, we work with specifications that yield expressions linear in parameters (for the log of marginal utility of consumption).

the marginal utility of consumption – which is mismeasured as a result. In this paper, we use the detailed data on both car purchases and ownership that are available in the CEX to condition on the consumption derived from ownership of durable goods. In order to derive a suitable specification by which to do this, we assume a utility function in which durable and non-durable consumption are non-separable, and durable consumption enters as an exponential – for example,

$$(16) \qquad U(c, K, z) = \frac{c^{1-\gamma}}{1-\gamma} \exp(\vartheta K + \delta' z),$$

where K is the stock of durables (cars) and z represents other taste shifters. This approach allows us to carry out a test of consumption insurance with non-separability between durables and non-durables using a similar specification to that used in deriving equations 11 and 12 but controlling for durable consumption (proxied by car consumption) in levels:

$$(11') \qquad \ln(c_{i,t+k}) = \varphi_i + \tilde{\delta}' z_{i,t+k} + \vartheta K_{i,t+1} + \tilde{\mu}_t + \lambda \ln(y_{i,t+k}) + \tilde{\omega}_{i,t+k}$$

$$(12') \qquad \Delta_k \ln(c_{i,t+k}) = \kappa + \tilde{\delta}' \Delta_k z_{i,t+k} + \vartheta \Delta_k K_{i,t+1} + \tilde{\lambda} \Delta_k \ln(y_{i,t+k}) + \Delta_k \tilde{\omega}_{i,t+k}.$$

Of course, there is a distinction between stocks and flows of durable goods that does not exist for non-durable goods. In particular, whilst the flow of non-durable goods is broadly equivalent to the consumption of those goods, it is the stock of durable goods that is likely to be more informative regarding consumption of the services provided by those goods. Thus our primary robustness check uses data on households' stock of cars. Such a test would be consistent with specifications 11′ and 12′.

To estimate equations 11′ and 12′, we will need measures of the stock of cars held by each household in the sample. Although we do have such information for a considerably long period, our data do not cover the entire period we analyse for non-durable consumption. In order to replicate our main analysis over the full time frame, we also carry out a robustness check that conditions on households' car purchases. While we will be considering group averages (by year-of-birth cohort and education), a structural interpretation of the results we obtain by adding the average purchase of cars to equations such as 11 and 12 is difficult. However, the presence of adjustment costs, and more generally the saliency of cars, makes expenditure on them particularly sensitive to income shocks, especially permanent or long-lasting shocks. Controlling for such a variable can therefore be quite informative.

3. Risk-sharing groups and aggregation

Until now, we have not specified the risk-sharing group for which we consider the social planner problem sketched above. In some situations, such as that considered in Townsend (1994), a risk-sharing group is naturally defined in the context one is considering. It could be a village (as in Townsend (1994)) or an extended family (as in Hayashi (1996) and Attanasio, Meghir and Mommaerts (2019)). And even if one states the implications of risk sharing for the whole economy, as an empirical strategy one can define a well-defined group and aggregate equations 11 and 12 for such groups to fit the available data. In particular, consider group g, with fixed and observable membership. Aggregating equations 11 and 12 over the members of such groups, one obtains

$$(17)\qquad \frac{1}{\#g}\sum_{i\in g}\ln\left(c^{g}_{i,t+k}\right)=\varphi^{g}_{i}+\tilde{\delta}'\frac{1}{\#g}\sum_{i\in g}z^{g}_{i,t+k}+\tilde{\mu}^{g}_{t}+\lambda\frac{1}{\#g}\sum_{i\in g}\ln\left(y^{g}_{i,t+k}\right)+\tilde{\omega}^{g}_{i,t+k}$$

$$(18)\qquad \frac{1}{\#g}\sum_{i\in g}\Delta_k\ln\left(c^{g}_{i,t+k}\right)=\kappa+\tilde{\delta}'\Delta_k\frac{1}{\#g}\sum_{i\in g}z^{g}_{i,t+k}+\tilde{\lambda}\Delta_k\frac{1}{\#g}\sum_{i\in g}\ln\left(y^{g}_{i,t+k}\right)+\Delta_k\tilde{\omega}^{g}_{i,t+k},$$

where $\#g$ is the size of group g and the sums are taken over all the members of group g. Equations 17 and 18 are essentially aggregations over group members of equations 11 and 12: the quantities that appear in them (consumption, income, demographics) are group means of the variables that appear in equations 11 and 12. For this reason, these equations can be extremely useful in the absence of longitudinal data. The group means in both equations can be estimated using time series of repeated cross sections, using the fact that the samples that make these cross sections can provide unbiased and consistent estimates of the means considered in equations 11 and 12. A similar argument applies for specifications that allow for other generalisations (such as equations 11′ and 12′), as long as they imply expressions for the log of marginal utility that are linear in parameters. For proper inference, the presence of this type of measurement error, which arises from the fact that the samples in the repeated cross sections are finite, has to be taken into account. The structure of the data allows the construction of appropriate instruments, as in Attanasio and Davis (1996). Furthermore, the computation of the standard errors for the estimated coefficients can take into account the presence of such a measurement error, which, in the case of equation 17, induces a moving average (MA) structure of the residuals of the equation. Below, we follow Attanasio and Davis in focusing on systematic relative wage movements across cohort–education groups of households.

This approach has several advantages. Unlike individual wage movements (which are likely to be correlated with individual preference shifts), systematic relative wage movements across large groups of workers are uncorrelated with idiosyncratic components of individual preference shifts. The focus on groups therefore allows us to deal with the presence of unobserved idiosyncratic shifts and measurement error. Furthermore, a focus on publicly observable relative wage movements means that evidence against consumption insurance cannot be explained by theories stressing the role of unobserved shocks in an informationally constrained optimal consumption allocation.[18] Finally, focusing on relative wage movements across groups of households allows us to make use of two (or more) rich cross-sectional data sets, one that contains high-quality information on consumption (which, however, lacks a longitudinal dimension[19]) and the other on earnings. This would not be possible without a single panel data set, containing information on both consumption and earnings, if our tests were carried out at the individual level.

While the focus on groups is very useful, the approach we use is not without problems. The biggest shortcoming lies in the fact that equations 17 and 18 test for the presence of risk sharing across groups. They are completely silent about the ability of insuring idiosyncratic shocks within groups. Indeed, purely idiosyncratic shocks are averaged within groups in equations 17 and 18 and are not considered at all.

V. Changes in relative wages, incomes and consumption

1. Testing consumption insurance

In this section, we detail the results of estimating the synthetic panel specifications derived in Section IV. We do so first using levels specifications, and then we report the results of difference specifications over various time frames.

Our baseline specification, estimated over cohort–education groups, contains the mean log pre-tax hourly wage among men, a polynomial in age, year effects and fixed effects for each cohort–education group for the levels specifications.

We also carry out estimation including household controls – which are the mean of log family size, the number of children under 3 years of age,

[18]For example, Green (1987), Townsend (1988), Phelan and Townsend (1991) and Atkeson and Lucas (1992).

[19]The CEX has a short panel dimension, in that the households in the sample report information for four consecutive quarters, but no households are followed for more than one year.

the number of other children and the number of adults. Doing so is intended as a way of capturing life-cycle preference variation that differs over time and across groups, which could be seen to drive a wedge between marginal utility growth and measured consumption growth. However, it is also plausible that to the extent that life-cycle consumption requirements do vary over time and across groups, this could itself be driven by variation in relative wages. Including these controls may thus make it less likely that we are able to reject consumption insurance.

In order to account for possible inconsistency of the ordinary least squares (OLS) estimator that could be caused by sampling variation or measurement error in our explanatory variables, we estimate specifications that instrument for wages and demographic variables. Since we use different data sources to construct the dependent and explanatory variables, sampling variation and measurement error in the explanatory variables will be uncorrelated with equation error in our regression model. We make the assumption that measurement error in each of our explanatory variables is uncorrelated with lags and leads, and with measurement error in other years. Our instruments in the levels specifications are thus constructed using a flexible combination of lags and leads, designed to avoid loss of observations due to instrumenting. Since measurement error and sampling variation in the levels induce a moving average error structure in differenced data, certain lags and leads are no longer valid instruments. For the one-year difference specification, we use an instrument that brackets the time interval of the true change (but does not involve the lagged measurement error); for multi-year specifications, we instrument with the lagged multi-year change.[20] Because this entails some loss of observations, the results reported below restrict estimation to a sample that is common across all specifications, given a particular differencing interval.

The first panel of Table 1 shows results for the levels specifications. Similarly to the Attanasio and Davis (1996) results for the period 1980 to 1990, we find a significant positive coefficient on male wages across all our levels specifications. In particular, the table shows a positive coefficient of 0.82 from OLS estimation not including demographic controls and a coefficient of 0.51 when demographic controls are included. When estimated instead using instruments (IV) as described above, both coefficients are higher – at 0.92 and 0.74 respectively – which is to be expected in the presence of measurement error. Full insurance against relative wage movements would imply coefficients equal to zero, so these estimates provide strong evidence against that hypothesis.

[20]Further details can be found on pages 1237–9 of Attanasio and Davis (1996).

TABLE 1
Synthetic panel regressions

	No controls OLS	*No controls IV*	*Controls OLS*	*Controls IV*
	Levels specifications ($N = 920$)			
Mean log wage	0.819***	0.920***	0.508***	0.738***
	(0.036)	(0.036)	(0.048)	(0.100)
	Annual difference specifications ($N = 764$)			
Mean log wage	0.108	0.295	0.070	0.235
	(0.082)	(0.183)	(0.071)	(0.195)
	Eight-year difference specifications ($N = 492$)			
Mean log wage	0.764***	0.877***	0.488***	0.840***
	(0.048)	(0.059)	(0.063)	(0.124)

Note: *** denotes that the effect is significantly different from zero at the 1 per cent level, ** at the 5 per cent level and * at the 10 per cent level. Standard errors are shown in parentheses. Household controls are the mean of log family size, the number of children under 3 years of age, the number of other children and the number of adults.

Next, Table 1 shows the results from difference specifications with an interval of one year. In contrast to the levels results, estimated coefficients across all four specifications are not significantly different from zero (although IV estimates are again larger). However, when the differencing interval is extended to eight years, estimated coefficients reflect those from the levels specifications – across all four estimates, they are statistically significantly different from zero, and the hypothesis of full insurance against relative wage movements is rejected.

Results from difference specifications with intervals of two and five years are shown in Table A3 in the online appendix. Both show estimated coefficients on the mean log wage that are significantly different from zero, with those from the five-year specification generally larger and significant at a higher level (the estimated coefficients from the difference specifications with an interval of eight years shown in Table 1 are larger still).

Taken together, these results seem to strongly reject the consumption insurance hypothesis. In fact, the results from both the eight-year difference specifications and the levels specifications are (excluding the eight-year OLS specification with demographic controls) closer to 1 than 0, suggesting substantial pass-through of relative wage movements into relative consumption growth. It is worth also noting, however, that the annual difference estimates suggest that the extreme alternative hypothesis of no consumption smoothing should also be disregarded.

2. Perfect insurance and non-separability with female labour supply

Tables 2 and 3 present a similar series of regressions, but add measures of female leisure and female wages respectively to the benchmark specifications. We instrument for these in a manner analogous to that described in the previous subsection. The coefficients on female leisure are imprecisely estimated in most specifications, but the OLS coefficients in both the levels and annual difference specifications are significantly negative whether or not household controls are included. These coefficients imply that increases in female leisure are associated with decreases in household consumption expenditures, suggesting preference non-separability in the form of a substitution relationship between leisure and consumption goods. Nevertheless, controlling for female leisure does not meaningfully impact on the size or statistical significance of coefficients on male wages in any of our specifications. Similarly, adding female wage measures as a control does not substantially change the size or significance of coefficients on male wages in any of our specifications.

TABLE 2

Synthetic panel regressions with female leisure

	No controls OLS	*No controls IV*	*Controls OLS*	*Controls IV*
	Levels specifications ($N = 920$)			
Mean log wage	0.819***	0.919***	0.508***	0.743***
	(0.036)	(0.036)	(0.047)	(0.101)
Female log leisure	–0.699**	–0.370	–0.491**	0.292
	(0.275)	(0.448)	(0.236)	(0.458)
	Annual difference specifications ($N = 764$)			
Mean log wage	0.118	0.288	0.079	0.240
	(0.082)	(0.185)	(0.071)	(0.196)
Female log leisure	–0.732**	0.522	–0.753**	–0.253
	(0.361)	(1.250)	(0.318)	(1.091)
	Eight-year difference specifications ($N = 492$)			
Mean log wage	0.777***	0.858***	0.762***	0.838***
	(0.048)	(0.061)	(0.047)	(0.209)
Female log leisure	–0.397	0.537	–0.349	0.209
	(0.355)	(0.542)	(0.347)	(0.505)

Note: *** denotes that the effect is significantly different from zero at the 1 per cent level, ** at the 5 per cent level and * at the 10 per cent level. Standard errors are shown in parentheses. Household controls are the mean of log family size, the number of children under 3 years of age, the number of other children and the number of adults.

TABLE 3

Synthetic panel regressions with female wages

	No controls OLS	*No controls IV*
	Levels specifications ($N = 920$)	
Mean log wage	0.707***	0.877***
	(0.053)	(0.068)
Female log wage	0.153***	0.053
	(0.052)	(0.069)
	Annual difference specifications ($N = 764$)	
Mean log wage	0.104	0.121
	(0.085)	(0.229)
Female log wage	0.013	0.294
	(0.067)	(0.211)
	Eight-year difference specifications ($N = 492$)	
Mean log wage	0.676***	0.819***
	(0.066)	(0.098)
Female log wage	0.130**	0.074
	(0.062)	(0.091)

Note: *** denotes that the effect is significantly different from zero at the 1 per cent level, ** at the 5 per cent level and * at the 10 per cent level. Standard errors are shown in parentheses. Household controls are the mean of log family size, the number of children under 3 years of age, the number of other children and the number of adults.

VI. Car stocks and expenditure: testing consumption insurance with non-separability between durables and non-durables

As discussed in Section IV, if preferences are non-separable between non-durable and durable consumption, the estimation above relies on a measure of the marginal utility of consumption that is mismeasured (the log of non-durable consumption). In this section, therefore, we augment our benchmark specifications with both the stock of cars and, separately, the expenditure on cars among each cohort–education group.

As shown in Table 4, adding in these measures changes the size of OLS estimates on male wages from the levels specifications only slightly, and they are still significantly different from zero, whether including household controls or not. Similarly, estimates from eight-year differencing intervals are broadly similar to those from the benchmark specifications, and all are still significantly different from zero. Annual difference specifications show no statistically significant estimates on male wages, again just as in the benchmark specifications. Similar observations can be made when comparing instrumental variable results, which are shown for specifications with car measures in

TABLE 4

Synthetic panel regressions with cars

	No controls OLS	*No controls OLS*	*Controls OLS*	*Controls OLS*
	Levels specifications ($N = 688$ [stocks]; 920 [expenditure])			
Mean log wage	0.667***	0.803***	0.348***	0.506***
	(0.039)	(0.036)	(0.049)	(0.047)
Stock of cars[a]	0.014***		0.013***	
	(0.001)		(0.001)	
Car expenditure[a]		0.017***		0.013***
		(0.003)		(0.002)
	Annual difference specifications ($N = 604$ [stocks]; 764 [expenditure])			
Mean log wage	0.057	0.099	0.000	0.066
	(0.090)	(0.081)	(0.075)	(0.070)
Stock of cars[a]	0.013***		0.012***	
	(0.001)		(0.001)	
Car expenditure[a]		0.013***		0.011***
		(0.002)		(0.002)
	Eight-year difference specifications ($N = 348$ [stocks]; 492 [expenditure])			
Mean log wage	0.689***	0.752***	0.411***	0.500***
	(0.055)	(0.047)	(0.065)	(0.062)
Stock of cars[a]	0.014***		0.014***	
	(0.002)		(0.001)	
Car expenditure[a]		0.020***		0.017***
		(0.003)		(0.003)

[a]In $ thousand.

Note: *** denotes that the effect is significantly different from zero at the 1 per cent level, ** at the 5 per cent level and * at the 10 per cent level. Standard errors are shown in parentheses. Household controls are the mean of log family size, the number of children under 3 years of age, the number of other children and the number of adults. Regressions including data on car stocks only cover years 1984 to 2010 due to data availability.

Table A4 in the online appendix. Thus, even when accounting for non-separability of preferences between durable and non-durable consumption, we find large and statistically significant departures from the consumption insurance hypothesis in both levels and eight-year differencing intervals, but not over annual differencing intervals. Our results also indicate that across the levels specifications and both differenced specifications, higher/increasing stocks of and expenditures on cars are associated with increases in household consumption non-durable expenditures, although the association is barely statistically significant.

VII. Conclusions

In this paper, we have observed, as has been established elsewhere, that inequality in family income rose substantially between 1980 and 2015, with the large part of this increase occurring in the early 1980s, the first half of the 1990s and since 2008. Relatedly, we have also noted that there have been pronounced changes in the structure of relative wages over the same period. We have shown, as Attanasio and Davis (1996) did for the 1980s, that relative wage movements among cohort–education groups of men drove large changes in the distribution of household consumption between 1980 and 2015. In particular, among households where the male head had less than 12 years of education, real household consumption fell through the 1980s and 1990s, particularly sharply in the early 1980s. This mirrors declines in real wages for most cohorts of these men across the same period. By contrast, real wages rose substantially across the period for college-educated men – reflected in their strong growth in real household consumption across the same period.

We have shown that Attanasio and Davis's rejection of the hypothesis of between-group insurance for non-durables across the 1980s holds when the estimation period is extended to cover 1980–2015. The estimated coefficients on male wages in the specifications derived are statistically significantly different from zero in both the levels specifications and in specifications of eight-year differences, indicating that relative wage changes do translate into relative consumption changes. Similarly to Attanasio and Davis, we find the coefficients are not significant in the annual difference specifications. These results hold across alternative specifications that include, variously, household demographics, female wages and female leisure as controls.

Furthermore, we extend this econometric analysis in order to overcome the potential issue of non-separability between durable and non-durable consumption. We do so by augmenting our benchmark specifications with car stocks and expenditures separately. Even after doing so, the consumption insurance hypothesis continues to be comprehensively rejected by our between-group analysis.

Supporting information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

- Appendix

References

Aguiar, M. and Bils, M. (2015), 'Has consumption inequality mirrored income inequality?', *American Economic Review*, vol. 105, pp. 2725–56.

Altonji, J., Hayashi, F. and Kotlikoff, L. (1992), 'Is the extended family altruistically linked? Direct tests using micro data', *American Economic Review*, vol. 82, pp. 1177–98.

Arellano, M., Blundell, R. and Bonhomme, S. (2017), 'Earnings and consumption dynamics: a nonlinear panel data framework', *Econometrica*, vol. 85, pp. 693–734.

Atkeson, A. and Lucas, R. (1992), 'On efficient distribution with private information', *Review of Economic Studies*, vol. 59, pp. 427–53.

Attanasio, O., Battistin, E. and Ichimura, H. (2007), 'What really happened to consumption inequality in the United States?', in E. E. Berndt and C. R. Hulten (eds), *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, Cambridge, MA: National Bureau of Economic Research.

— and Davis, S. J. (1996), 'Relative wage movements and the distribution of consumption', *Journal of Political Economy*, vol. 104, pp. 1227–62.

—, Hurst, E. and Pistaferri, L. (2014), 'The evolution of income, consumption and leisure inequality in the US, 1980–2010', in C. Carroll, T. Crossley and J. Sabelhaus (eds), *Improving the Measurement of Consumer Expenditures*, Chicago, IL: University of Chicago Press.

—, Meghir, C. and Mommaerts, C. (2019), 'Insurance in extended family networks', National Bureau of Economic Research (NBER), Working Paper no. 21059.

— and Pistaferri, L. (2016), 'Consumption inequality', *Journal of Economic Perspectives*, vol. 30, no. 2, pp. 3–28.

Bar-Ilan, A. and Blinder, A. (1992), 'Consumer durables: evidence on the optimality of usually doing nothing', *Journal of Money, Credit and Banking*, vol. 24, pp. 258–72.

Betson, D. M. (2004), 'Poverty equivalence scales: adjustment for demographic differences across families', paper presented at the National Research Council Workshop on Experimental Poverty Measures, Washington DC, 15–16 June; accessed at https://pdfs.semanticscholar.org/5944/790667489b5a8189bc0c680cdf5ae2c51701.pdf.

Blundell, R., Pistaferri, L. and Preston, I. (2008), 'Consumption inequality and partial insurance', *American Economic Review*, vol. 98, pp. 1887–921.

—, — and Saporta-Eksten, I. (2016), 'Consumption inequality and family labor supply', *American Economic Review*, vol. 106, pp. 387–435.

—, — and — (2018), 'Children, time allocation, and consumption insurance', *Journal of Political Economy*, vol. 126, pp. S73–115.

— and Preston, I. (1998), 'Consumption inequality and income uncertainty', *Quarterly Journal of Economics*, vol. 103, pp. 603–40.

Cochrane, J. (1991), 'A simple test of consumption insurance', *Journal of Political Economy*, vol. 99, pp. 957–76.

Cutler, D. and Katz, L. (1991), 'Macroeconomic performance and the disadvantaged', *Brookings Papers on Economic Activity*, vol. 1991, no. 2, pp. 1–61.

Dynarski, S. and Gruber, J. (1997), 'Can families smooth variable earnings?', *Brookings Papers on Economic Activity*, vol. 1997, no. 1, pp. 229–303.

Green, E. (1987), 'Lending and the smoothing of uninsurable income', in E. Prescott and N. Wallace (eds), *Contractual Arrangements for Intertemporal Trade*, Minneapolis, MN: University of Minnesota Press.

Greenspan, A. and Cohen, D. (1996), 'Motor vehicle stocks, scrappage and sales', Board of Governors of the Federal Reserve System, Finance and Economic Discussion Paper no. 96-40.

Hayashi, F. (1996), 'Analysis of household saving: past, present, and future', *Japanese Economic Review*, vol. 47, pp. 21–33.

Krueger, D. and Perri, F. (2006), 'Does income inequality lead to consumption inequality? Evidence and theory', *Review of Economic Studies*, vol. 73, pp. 163–93.

Leamer, E. (2007), ‘Housing is the business cycle’, National Bureau of Economic Research (NBER), Working Paper no. 13428.

Mace, B. (1991), ‘Full insurance in the presence of aggregate uncertainty’, *Journal of Political Economy*, vol. 99, pp. 928–56.

McCarthy, M., Johnson, D., Garner, T. and Passero, B. (2002), ‘Issues in construction and research use of the Consumer Expenditure Survey’, paper presented at the 2002 NBER Summer Institute, Cambridge, MA.

Meyer, B. and Sullivan, J. (2010), ‘Five decades of consumption and income poverty’, Milton Friedman Institute (MFI), Working Paper no. 2010-003.

— and — (2017), ‘Consumption and income inequality in the U.S. since the 1960s’, National Bureau of Economic Research (NBER), Working Paper no. 23655.

Phelan, C. and Townsend, R. (1991), ‘Computing multi-period, information-constrained optima’, *Review of Economic Studies*, vol. 58, pp. 853–91.

Townsend, R. (1988), ‘Information constrained insurance: the revelation principle extended’, *Journal of Monetary Economics*, vol. 21, pp. 411–50.

— (1994), ‘Risk and insurance in village India’, *Econometrica*, vol. 62, pp. 539–91.