Time Variation in Lifecycle Consumption and Income*

Yunus Aksoy
Birkbeck, U of London

Henrique S. Basso
Banco de España

Carolyn St Aubyn
Birkbeck, U of London

2nd November 2021

Abstract

We document systematic and significant time variation in US lifecycle non-durable consumption profiles. Consumption profiles have consistently become flatter: intergenerational differences in consumption across age groups have decreased over time. Pooling data across different periods to identify lifecycle profiles and failing to account for unobserved heterogeneity masks relevant time variations and may artificially generate hump-shaped consumption age profiles. The main driver behind lifecycle consumption variations are lifecycle income changes, which display similar flattening. Employing a lifecycle model we show changes in income are sufficient to match the movements in consumption.

Key Words: Age Profile of Consumption, Age Profile of Income, Consumption Heterogeneity, Pooling

JEL Classification Codes: E21, J11

---

*Aksoy and St Aubyn: Birkbeck, University of London, Malet Street, WC1E 7HX, London, United Kingdom, e-mails: y.aksoy@bbk.ac.uk, c.molesworthstaubyn@bbk.ac.uk. Basso: Banco de España, DG Economics, Statistics and Research, Alcalá 48, 28014, Madrid, Spain. e-mail: henrique.basso@bde.es. We first would like to thank, without implicating, Ron P. Smith for his invaluable advice. We thank Zühre Aksoy, Mehmet Barlo, Walter Beckert, Jesús Fernández-Villaverde, Alpay Filiztekin, Nezih Güner, Andrew Hugh Hallett (discussant), Özgür Kıbrıs, Yuliya Kulikova, John Muellbauer, Ernesto Villanueva, Ünal Zenginobuz and Gylfi Zoëga and seminar participants at Birkbeck, Boğaziçi, Sabancı, Surrey, Banco de España, CBS/CBI Research Seminar in Reykjavik, CEF in Ottawa, MMF LSE conference, Durham Macroeconomic Policy Workshop, Oxford NuCamp workshop for comments. Aksoy acknowledges a Birkbeck BEI small research grant and St Aubyn PhD financial support from the ESRC (ES/J500021/1). Part of this research was conducted when Aksoy was visiting ADG Economics and Research, Banco de España, as a Resident Visiting Scholar. Aksoy and Basso are affiliated with Birkbeck Centre for Applied Macroeconomics (BCAM). The views expressed in this paper are those of the authors and do not necessarily coincide with those of the Banco de España and the Eurosystem.

1 Introduction

The lifecycle profile of nondurable consumption expenditures, defined as the curve that depicts the level of nondurable consumption expenditures across ages, has been studied in the seminal papers of Deaton and Paxson (1994), Attanasio et al. (1999) and Gourinchas and Parker (2002), and more recently by Fernandez-Villaverde and Krueger (2007) and Aguiar and Hurst (2013). The consensus view is that consumption expenditures increase with income in the earlier part of the lifecycle, are hump-shaped, peaking around the age of 55 and falling at the later part of the lifecycle.

When analysing lifecycle patterns of consumption, a commonly made implicit assumption is that across time (waves) households of the same age behave in a similar fashion and face similar age specific structural economic conditions. Data is thus pooled across time. This approach may be misleading, particularly when household unobserved heterogeneity is not accounted for. Given changes in macroeconomic and microeconomic conditions, in the technological environment and mode of production, in demographic structures and in the evolution of asset prices and income, the homogeneity assumption in consumption decisions of households of a given age across waves needs validation and cannot be taken at face value. Hence, we relax this assumption and after controlling for household characteristics, we show that US lifecycle nondurable consumption profiles have consistently become flatter over time, indicating that at any given point in time, intergenerational consumption differences have declined.

Lifecycle nondurable consumption expenditures and income can be viewed from an age (Figure 1) or from a cohort (Figure 2) perspective. It is clear that Figures 1 and 2 depict an observationally equivalent hump-shaped consumption expenditures and income lifecycle profiles peaking roughly around the age of 55. This highlights the fact that one cannot distinguish between an explanation which says that, after time fixed effects affecting all households are controlled for, a household consumption is determined by household head's *age*, from an explanation which says it is determined by household head's year of birth, that is $Cohort = Year - Age$.

Posing a question of whether we observe systematic difference in consumption patterns across ages rather than cohorts is more productive for several reasons. First, it is parsimonious. Given the life span, there are a fixed number of ages whereas the number of cohorts keeps increasing. Second, there is extensive economic theory about the lifecycle, but relatively little about the patterns of cohort effects. Finally, intergenerational differences can be interpreted more directly when we focus on life-cycle profiles presented from an age perspective. Thus, we primarily take a age based view to analyse lifecycle consumption and later study the role of cohort effects on our results.

Figure 1: Unconditional Nondurable Consumption Expenditures and Income by Age
Note: Top curve shows the income for different age groups. Bottom curve shows the nondurable consumption expenditures for these age groups.
Source: PSID 1998-2014

We study consumption expenditures using a longitudinal panel of US households

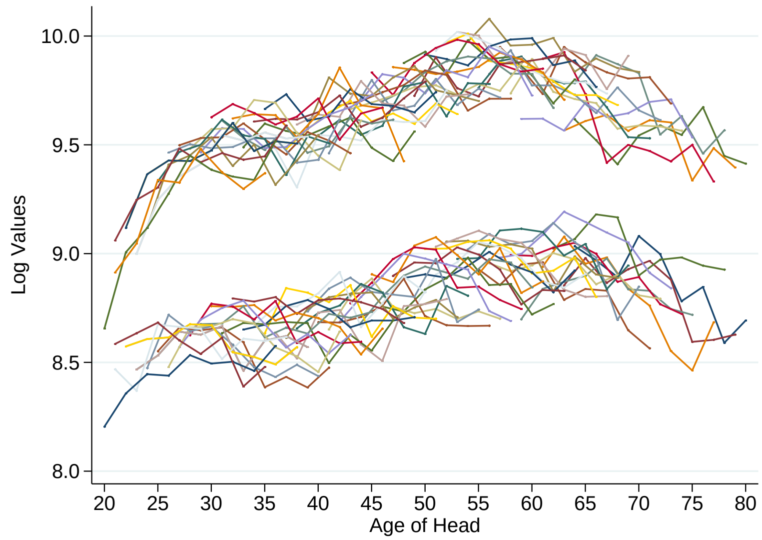

Figure 2: Unconditional Nondurable Consumption Expenditures and Income by Cohort
Note:Top curves shows the income for different cohorts. Bottom curve shows the nondurable consumption expenditures for these cohorts.
Source: PSID 1998-2014

for the period 1998-2014. Our identification strategy relies on the aggregation of households into age groups of 5 and 15 years rather than including directly the household's age. For each wave, the time fixed effects vary, some households move to the next age group while others remain, providing the necessary variation to identify time varying *age group* effects after controlling for time and household fixed effects. First by pooling all the data and ignoring age-group time variation we confirm nondurable consumption expenditures display lifecycle properties and are hump-shaped in line with the literature (e.g. Attanasio et al. (1999) and Aguiar and Hurst (2013)). We then allow for age-group and time interactions and document that there are systematic and significant time variations in lifecycle consumption expenditures of households in the US in the sample period we study. We show that differences between consumption expenditures across age groups have declined and lifecycle or age consumption profiles became flatter over time. Furthermore, we find that age-group specific parameters of older households display more variation over time, and observe that for none of the years (waves) lifecycle profiles are hump-shaped. Thus, pooling data across different periods and not controlling for household fixed effects to identify consumption profiles introduce estimation bias, mask significant and economically relevant time variation, and may artificially generate the well known hump-shaped lifecycle consumption profiles.

These results are robust with respect to longer longitudinal data covering the period 1980-2014, altering the size of age groups (necessary to identify age-group effects when time and age-group fixed effect interactions are included), education levels, the inclusion of household level economic controls (income and housing wealth), the exclusion of households who do not own a house, and different ways of adjusting for family size. We also verify that age profiles are not affected when cohort effects are controlled for. Finally, while PSID data does not have complete lifecycle data for individual cohorts, for the sake of completeness, we also present results for time-varying cohort effects. Although the empirical model is more limited to ensure cohort groups are stable and well populated across the sample, our time-varying cohort estimations show similar patterns as the time-varying lifecycles based on age.

Estimating the model with a standard OLS specification prevents the inclusion of controls for the pervasive household specific unobserved heterogeneities and therefore introduces substantial bias in the identification of lifecycle profiles as, for instance, in the case when the Consumer Expenditure Survey (CEX) is used. The PSID panel data allows us to estimate the model with a fixed effects specification and thus we address this issue, accounting for unobserved heterogeneity in our methodology. Nonetheless, when we estimate an OLS using CEX and PSID data we obtain similar age-group profiles, indicating that the results are not driven by the potential differences between the CEX and PSID datasets. Aguiar and Hurst (2013), using food data from the

PSID, obtain similar results estimating an OLS and a fixed effect model, inferring that biases from the lack of fixed effects may not be relevant in this setting. We find their conclusion cannot be extended to nondurable consumption. In the case of nondurable consumption, controlling for household unobserved heterogeneity is crucial to obtain unbiased empirical estimates, particularly when age-group and time interactions are considered.

Aguiar and Hurst (2013) also study the lifecycle consumption of different expenditure categories and find that work-related consumption expenditures, such as clothing, transportation and food away from home decline as households get older, driving the hump-shaped nature of consumption profiles. We observe a flattening of lifecycle consumption profiles through time in almost all sub-categories in our sample, including work-related categories, such as transportation and food-away, suggesting our conclusions extend to consumption sub-categories.

The key intuition motivating our analysis is that, when studying lifecycle consumption, agents of the same age-group at different points in time should not be treated as homogeneous and, as a consequence, should not be pooled together. That is subtly distinct from taking into account cohort effects only. In fact, age-group profiles are robust to the inclusion of the households' birth year as an additional control, confirming constant cohort effects are not driving our results. We document that with the systematic flattening of lifecycle consumption profiles the difference of consumption across cohorts varies through time. The appropriate interpretation therefore is not that we identify constant cohort effects but rather that cohort effects are systematically changing through time. These results do not imply that 35 year old households today are relatively better off than 35 year old households in the 1990s, or that consumption inequality has changed through time, rather, the results indicate that at each fixed point in time throughout the period we study, after controlling for household fixed effects, the intergenerational consumption differences have decreased. Our results, together with findings of the extensive literature that document widening of consumption and income inequalities since 1980's in most advanced economies (see for instance Aguiar and Bils (2015) and Hoffmann et al. (2020) and references therein), indicate that through time the increased inequality may be more likely explained by household specific characteristics and not by age-group or lifecycle effects.

Finally, the time variation in age-group effects may reflect the differential effects of the business cycles on each age group. However, if that were to be the case, we should observe a substantial shift in the lifecycle profiles during the great recession (2008), in a similar fashion to the changes we observed in the estimated time fixed effects. Instead, our evidence points to slow moving and more systematic shifts in lifecycle profiles that occur throughout our sample period, indicating the more plausible interpretation is that the relationship between consumption levels across age groups has been structurally changing in recent decades.

What might be behind this time variation in the consumption profiles we uncover? Gourinchas and Parker (2002) stress the importance of the expected growth rate of income in determining consumption behaviour as households age and Attanasio et al. (1999) find that groups of households characterised by a relatively steeper income profile also present a steeper consumption profile, indicating the evolution of income in the lifecycle is a key driver of age-group consumption profiles. After controlling for the age-group specific component that depends on the lifecycle income, we find that consumption profiles are no longer flattening. Higher income in the lifecycle has become strongly associated with higher consumption levels. To confirm the relevance of time variation in income profiles as the driver of our results, we extract the age-group specific profiles of income following the same procedure as the one applied to consumption. In line with the work of Kambourov and Manovskii (2009) and Jeong et al. (2015), who report changes in lifecycle earnings due to a fall in the price of experience, we find very similar patterns of time variation in income to the one we observe for consumption: income lifecycle profiles have also become systematically flatter.[1] We perform the same robustness exercises for income, as done by consumption and find that the systematic time variation in income persist in all cases. Finally, we investigate whether housing wealth may be also driving variations in lifecycle consumption and find it not to be the case. Although, in line with the literature, we find that changes in the subjective housing wealth significantly affect consumption, particularly for older households, controlling for housing does not qualitatively alter the lifecycle consumption flattening we observe.

In our last empirical exercise we estimate consumption and income lifecycle profiles with panel data from the Italian Survey of Household Income and Wealth (SHIW) and find similar results: age-group patterns for both consumption and income have been flattening and accounting for unobserved heterogeneity alters results substantially. Thus, the systematic time variation uncovered is not restricted to the US, but may be a more general feature.

Finally, we provide a theoretical justification to our empirical findings. By employing a lifecycle model with consumption, housing and liquid assets choices, and feeding the estimated changes in the age-group profile of income, we find that variations in lifecycle income are sufficient to generate the observed changes in consumption profiles, confirming our suggestive evidence that income is driving the systematic changes

[1] Elsby and Shapiro (2012) also report a decrease in the experience earnings profiles for low skilled workers.

in the age-group profiles of consumption. We find that incorporating higher volume of credit and house price changes do not significantly affect consumption profiles. However, only after incorporating the changes in credit markets and the dynamics of house prices observed from 1998 onwards, the match between the theoretical lifecycle patterns of asset holdings (housing and non-housing wealth) and the one observed in the data improves.

The remainder of the paper is organised as follows. Data, the econometric methodology and results are presented in Section 2. In Section 2.2 we present our benchmark time-varying lifecycle results, and Section 2.3 investigates whether subjective house valuation and income are behind the variation in lifecycle consumption profiles we observe. The theoretical model is discussed in Section 3. Finally, Section 4 concludes.

2 Empirical Evidence

We study nondurable consumption expenditures using a longitudinal panel of US households that allows us to determine the age-group effects after controlling for household characteristics, fixed effects, income and perceived housing wealth. We then consider whether age-group effects depend on housing wealth and income. We start by presenting the data, then discuss the methodology and main empirical results.

2.1 Data

Data are from the nationally representative longitudinal US household survey, the Panel Study for Income Dynamics (PSID).[2] The survey was conducted annually from

[2] An alternative data set is the CEX. In general, this is considered the gold standard of consumption data in the US. The PSID is selected over the CEX because of its longitudinal structure. This

1968 to 1997[3] and biannually thereafter. It contains detailed information on household employment, income, consumption, assets and various household characteristics such as health status and social behaviour of around 5000 households (about 18,000 individuals) and their descendants with the addition of new households to maintain a nationally representative sample.[4]

Nondurable consumption expenditures, $c_{i,t}$, is defined as the sum of imputed rent, house insurance, utilities, nondurable vehicle costs, childcare, education costs, health insurance, nondurable transport costs such as parking, cabs and public transport, medical expenses, food at home, food away from home and the cash value of food stamps.[5]

The benchmark sample, using data from 1998 up until 2014, capitalises on the expanded nondurable consumption questions introduced in 1999 (data labelled 1998). This additional information, listed above, is used to construct a full measure of nondurable consumption. We have 42,720 observations. The average length of household participation in the survey in this data set is 6.7 waves, with a maximum of 9 waves (40.45 percent) and a minimum of one wave (3.3 percent). About 66.28% of house-

allows us to control for unobserved household effects which is not possible in the CEX. Also, we capitalise on the expanded consumption questions introduced in the PSID in 1999. With this, the consumption in the PSID covers 70% of the consumption measured in the CEX (Li et al. (2010)). We show that OLS estimations using CEX or PSID data deliver similar lifecycle profiles.

[3]Each wave of the survey asks households about the previous year's expenditures. We follow convention by labelling each wave, t as time period $t-1$. This means that information gathered in the 2003 wave will be labelled in the data set as 2002.

[4]For a full explanation of sample selection see Appendix A.

[5]As is standard in the literature, these expenditures act as a proxy for consumption. In fact, it underestimates the true amount by not accounting for consumption of leisure, home production and durable goods but assumes separable utility between these groups. Estimating the age-group profile over different categories; total consumption expenditures, nondurables and durables all yield the hump shape over the lifecycle (Fernandez-Villaverde and Krueger (2007)). Mankiw (1982) points out that durables and nondurables differ only in their rate of depreciation and that some nondurables, for example, clothing, are partly durable. So if the weight of durability relates to the type of consumption then the mix matters. Also, simply removing perceived durables is not sufficient to exclude durability altogether.

holds in the sample are homeowners. For robustness we repeat our analysis over a longer time period 1980 to 2014 based on imputed data as in Attanasio and Pistaferri (2014). We also report robustness analysis based on different methods of deflating the consumption data as in Aguiar and Hurst (2013) and how best to adjust for household size and composition. (see Section 2.2.3)

In some specifications we include a measure of total family income. The PSID includes a number of measures of income and earnings. We define total family money income $Y_{i,t}$ as the sum of taxable family income, family transfers and social security benefits. Taxable family money income is the sum of the head's asset income (dividends, interest, rental income and asset income from farm business), the spouse's asset income, and head and spouse labour income. Family transfer income consists of transfer income for family members other than husband and wife and aid to dependent children. All income measures are deflated and scaled following the same procedure adopted for consumption.

Finally, in some specifications we include a measure of subjective housing wealth. Our preferred subjective home value proxy is based on the responses of homeowners to a question in the PSID survey and reported in housing, mortgage distress and wealth data. Specifically homeowners are asked:

> 'A20. Could you tell me what the present value of (your/their) (apartment/mobile home/house) is (including the value of the lot if (you/they) own the lot)–I mean about how much would it bring if (you/they) sold it today?'

The question offers an insight into subjective expectations of households about their perceived wealth over a 50 year time period. Household responses to this question define our subjective variable $H_{i,t}$=Subjective Current Home Value. The average values of $H_{i,t}$ in our sample are strongly correlated (correlation coefficient: 0.96) with the Case-Shiller House Price Index.[6]

2.2 Lifecycle Consumption Profiles

The standard approach (e.g. Aguiar and Hurst (2013)) to estimate nondurable consumption expenditures utilising repeated cross sectional surveys, such the Consumer Expenditure Survey (CEX), relies on a specification such as:

$$c_{i,t} = \alpha + \beta_{age} Age_{i,t} + \delta_t D_{i,t}^{Time} + \zeta_C D_i^{Cohort} + \psi_Z Z_{i,t} + \epsilon_{i,t} \qquad \textit{Literature Benchmark} \quad (1)$$

where $c_{i,t}$ is nondurable consumption of household i during year t on consumption, $Age_{i,t}$ is a vector of year age dummies referring to the age of the household head, $D_{i,t}^{Cohort}$ is a vector of 1-year birth cohort dummies, D_t is a vector of year dummies, and $Z_{i,t}$ is a vector of family characteristics. As it is well known, in this setting, age, year, and cohort effects are not identified, because $Age + Cohort = Year$.

In order to identify and isolate the age component of consumption, or age profiles, Aguiar and Hurst (2013) employs the approach implemented by Deaton and Paxson (1994), attributing consumption growth to age and cohort effects and using normalised year dummies to capture cyclical fluctuations (i.e. orthogonalised to a time trend and add to zero). Another alternative is to include cohort effects but similarly assume they are orthogonal to a time trend (for further discussion see Schulhofer-Wohl (2018)). Aguiar and Hurst (2013), p.449, report that a more parsimonious represent-

[6]We also consider a proxy for subjective housing equity (SHE, house value net of outstanding mortgages). Inclusion of SHE does not improve our results and are available from authors.

ation that incorporates only time and age effects delivers similar age profiles. As such, we select as our benchmark the parsimonious model without cohort effects and later confirm, similarly to Aguiar and Hurst (2013), that our results are not altered when cohort effects are included employing both standard approaches, when we impose an additional restriction by normalising either time or cohort effects.

Unlike a repeated cross sectional survey, like CEX, PSID is a panel survey tracking individual households over time. Thus our estimation strategy leverages this panel dimension. We estimate the model with unbiased panel fixed effects estimator rather than the OLS and explicitly account for unobserved heterogeneity that is pervasive in household panel surveys.[7] We postulate that the log of nondurable consumption expenditures $c_{i,t}$ for each household $i = 1, ..., N$ at time $t = 1998, ..., 2014$ depends on a households fixed effects α_i, on a set of time-varying household characteristics,[8] $Z_{i,t}$, a vector of year dummies $D_{i,t}^{Time}$ and on age-group effects described by a group of dummies denoted $Age_{i,g,t}$, to capture lifecycle patterns.

We consider two models for the age-group dependent control $Age_{i,g,t}$. In the first model, denoted *Pooled Lifecycle*, and in line with the literature (e.g. Aguiar and Hurst (2013)), we assume lifecycle effects do not change over time, setting $Age_{i,g,t} = Age_{i,g} = D_{i,g}^{Age}$, where $D_{i,g}^{Age}$ is a dummy variable that takes the value one if the age of the head of household i is within the age group g and zero otherwise. $\beta_{g,t} = \beta_g$ in this case captures the log difference in mean consumption of the youngest age group (reference group) to the other age groups for the entire sample period. Formally, the

[7] Advantages of estimating with panel fixed effects against the OLS are well-known. In Section 2.2.2 and Appendix Appendix B.6 we provide further discussion and detailed comparisons of alternative econometric models in terms of their fit.

[8] These include dummy variables for the level education of the head of the household (grade school only, high school education, incomplete university education, and a university degree or higher), dummy variables for the number of children and adults in the household, race, marital status, state of residence and home ownership.

benchmark fixed effects specification is

$$c_{i,t} = \alpha_i + \beta_g D_{i,g}^{Age} + \delta_t D_{i,t}^{Time} + \psi_Z Z_{i,t} + \epsilon_{i,t} \qquad \textit{Pooled Lifecycle} \quad (2)$$

First, as in the case of Aguiar and Hurst (2013) our specification ensures that the time effects capture the variation common to all households at each point in time. Second, the household fixed effects capture the specific variation for each household that is common across all the time periods, leveraging the panel dimension of the dataset. Finally the age-group effects capture the remaining variation that is common for all households within the age group g. We initially select 4 fifteen-years age groups $g = 1, \ldots, 4 \quad (24-35, 36-50, 51-65, 65+)$ to ensure each age group is well populated but also consider 10 five-years age groups, with $g = 1, \ldots, 10 \quad (24-30, ..65-70, 71+)$ for robustness. By considering age groups of 5 and 15 years we avoid the traditional problem that age and time would move in tandem and identification would not be possible. In our setting, for each wave, the time fixed effect vary, some households move to the next age group while others remain, providing the necessary variation to identify age-group effects. In Section 2.2.1 below we augment the model to also include cohort effects for robustness.

Our second model, denoted *Time-varying Lifecycle*, accounts for time variation in lifecycle consumption expenditures by setting $Age_{i,g,t} = D_{i,g}^{Age} + (D_{i,g}^{Age} \times D_{i,t}^{Time}$, for $t \geqslant 2)$, thus adding an interaction term of age and time dummies. In this specification we allow the consumption allocations, not explained by household characteristics and business cycles effects, of an age group g to potentially change with time. $\gamma_{g,t}^1$ captures the log difference in mean consumption of age group g to the youngest age group (reference group) in the reference year (which we set to be the first wave, $t = 1$) and

$\gamma^2_{g,t}$ the added difference for each subsequent wave/year. Thus,

$$c_{i,t} = \alpha_i + \gamma^1_{g,t} D^{Age}_{i,g} + \gamma^2_{g,t} D^{Age}_{i,g} \times D^{Time}_{i,t} + \delta_t D^{Time}_{i,t} + \psi_Z Z_{i,t} + \epsilon_{i,t}. \tag{3}$$

Note that this model can be reparameterised, setting $\gamma_{g,t} = \gamma^1_{g,t} + \gamma^2_{g,t}$, obtaining simply

$$c_{i,t} = \alpha_i + \gamma_{g,t} D^{Age}_{i,g} \times D^{Time}_{i,t} + \delta_t D^{Time}_{i,t} + \psi_Z Z_{i,t} + \epsilon_{i,t} \quad \textit{Time-varying Lifecycle}. \tag{4}$$

Equation 4 allows for a direct comparison between the pooled estimate β_g and the time-varying age effects $\gamma_{g,t}$ for each t, which in this specification captures the log difference in mean consumption of age group g to the youngest age group (reference group) for each wave/year in our sample, providing a full representation of age-group specific effects on consumption for each year.

In Figure 3 we display the lifecycle coefficients (β_gs and $\gamma_{g,t}$s) or the age-group effects for each specification.[9] The upper panel shows estimates for broadly defined (15 years) and lower panel for narrowly (5 years) defined age groups. The thick dark line shows the coefficients from the regression that pools the information over the entire sample to measure age-group effects (β_g). The results are well known and depict a hump-shaped pattern of consumption in the lifecycle. The dashed lines display the lifecycle coefficients ($\gamma_{g,t}$) for each year $(1998, \ldots, 2014)$ separately.[10] We observe a systematic time variation in lifecycle consumption patterns. At the beginning of the sample (1998 - 2000) consumption is increasing in age groups. With time the lifecycle

[9] Table A.5 in the Appendix shows estimation results for i. the Pooled Lifecycle model, ii. the Time-varying Lifecycle model (Benchmark) *iii.* Time-varying Lifecycle model with economic controls including total family income and subjective house value as additional controls.

[10] Although qualitatively comparable, β_g's for the pooled age-group effects regression (thick dark line) are not a simple first order function of the $\gamma_{g,t}$'s estimated for each year/wave (dashed lines).

profile pivots down and towards the end of the sample period (2014) consumption is decreasing in age groups.

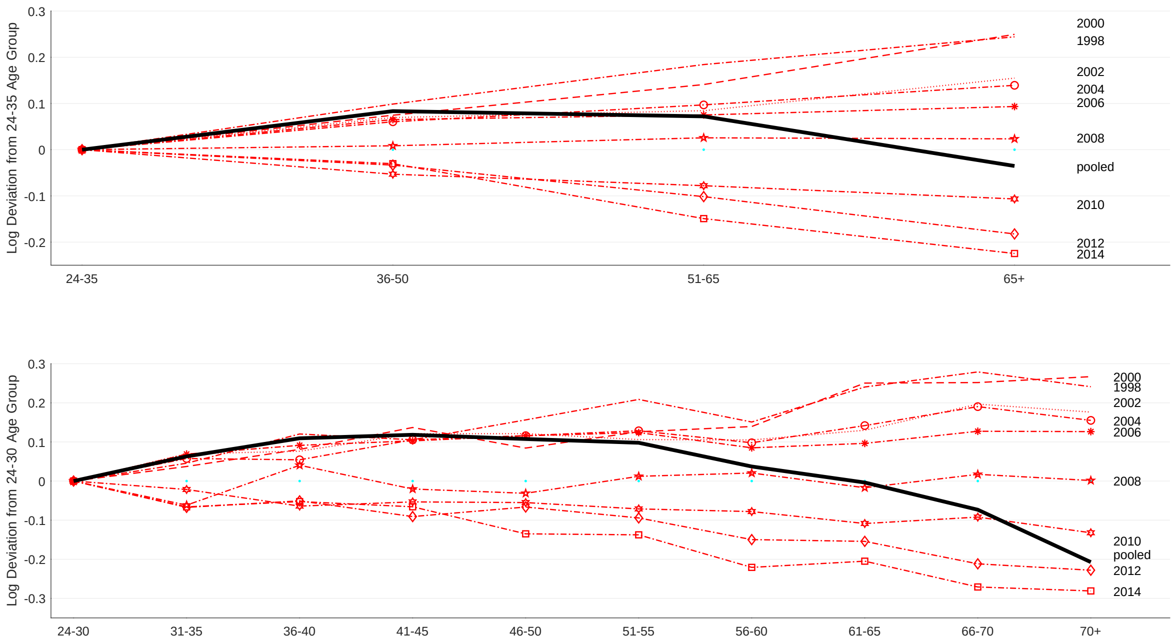

Figure 3: Lifecycle Consumption Patterns

Note: Each dashed line depicts $\gamma_{g,t}$ for each year of the wave of the survey, $(1998, \ldots, 2014)$, describing the estimated lifecycle pattern of consumption for each year. The dark line depicts the age-group effects β_g which pools information for the entire sample. The top graph considers 4 age groups, while the bottom graph show the results for 10 age groups.

In Figure 4 we show the age-group effect coefficients and their 90% confidence interval organised in the form of a time series (in the top panel). We plot $\gamma_{g,t}$ by age group over all time periods, comparing the within age group changes across time. To assess the significance of these changes, we test the hypothesis that the coefficients for each age group do not change over time, formally $H : \gamma_{g=s,t=1998} = \gamma_{g=s,t=1998+x}$, which is an implicit assumption of the pooled approach. Table 1 shows that the null hypothesis is rejected in almost all cases. We observe that age group coefficients, showing the relative difference w.r.t. the young age group (24-35) for each year, are economically and statistically different from each other. In Figure 4 (bottom panel) we also present age group coefficients grouped by time. Set out this way, these represent a sequence of lifecycle consumption profiles. This further illustrates the decrease in slope of lifecycle consumption profiles, with older age groups experiencing larger consumption variations than the young and middle aged households. We also compare the information criteria for these two specifications, *Pooled Lifecycle* versus *Time-varying Lifecycle*. Because the latter nests the former, we can use the information criteria as a likelihood ratio test with a penalty for complexity. Two popular information criteria, AIC and BIC, favour time variation in age-group effects. We also apply this test to the more granular age group specification and find strong evidence that allowing age-group effects to vary in time fits the data significantly better than pooling age-group effects over time (see Table A.5 in the Appendix for details).

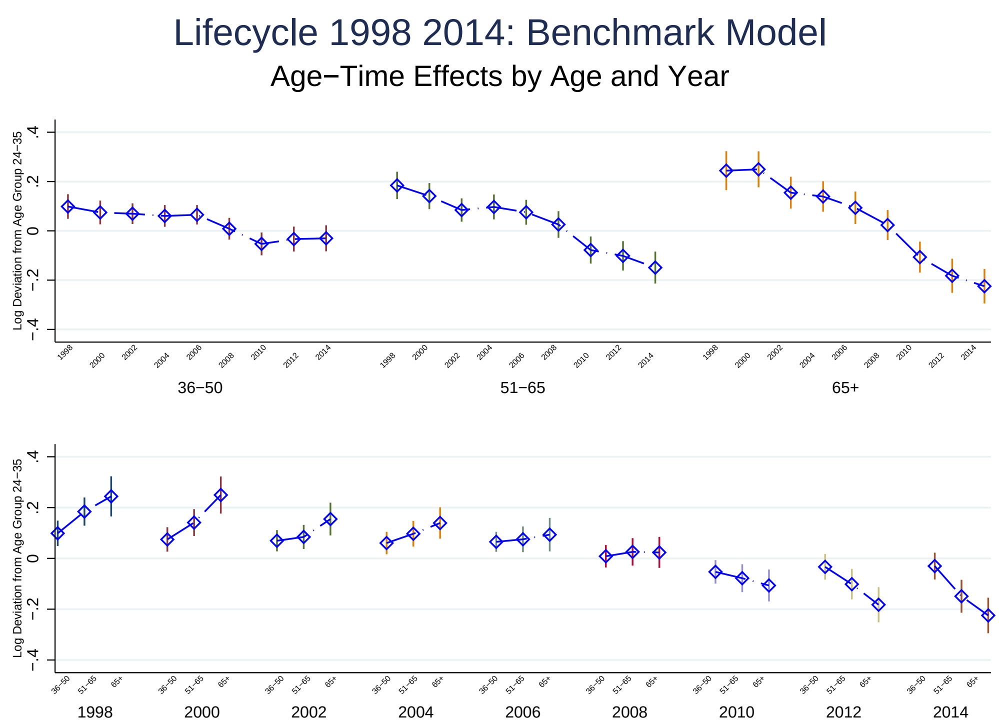

Figure 4: Age group coefficients ($\gamma_{g,t}$) plotted by age group (Top) and by year (Bottom).

Finally, we complement our lifecycle estimates for each wave by taking into account the consumption behaviour of the reference (youngest age-group) group (24-35), which may be changing over time. In order to evaluate this potential business cycle effects

Age Group	Year 2002	2006	2010	2014
35-50	0.2253	0.2871	0.0000	0.0062
51-65	0.0009	0.0097	0.0000	0.0000
>65	0.0280	0.0014	0.0000	0.0000

Table 1: $\gamma_{g,t}$ - Time Variation Statistical Test

Note: We test the hypothesis that the coefficients for each age group do not change over time. Results are shown for the base year, 1998 against 2002, 2006, 2010 and 2014, $\gamma_{g=s,t=1998} = \gamma_{g=s,t=1998+x}$.

on the evolution of the lifecycle of the reference group we re-estimate the model where the reference group now is 24-35 age group in 1998 (thus we drop time dummies to avoid perfect collinearity). Figure 5 records the coefficient estimates for the young age group w.r.t. the 1998 reference year.[11] Consumption expenditures of the young age group have in general drifted up from 1998 till 2014 (with a large fall and subsequent recovery due to the 2008-9 crisis).

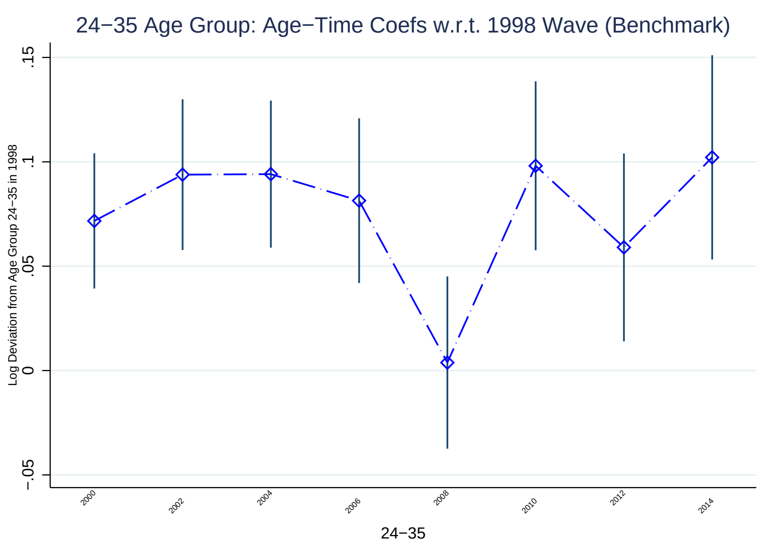

Figure 5: Consumption Drifts in Age Group 24-35

The evidence presented in Figure 5 together with Figure 4 suggests that lifecycle consumption expenditure evolution has at least two dimensions. One is related to the shifts in the consumption behaviour of the young age group with respect to the

[11]This exercise is equivalent to reporting the time dummies δ_t's in Equation 4.

business cycle; the other relates to the systematic shifts in the consumption behaviour of older age groups with respect to the young. Another potential interpretation is that the interaction term between age-group and time reflects the differential effects of the business cycles on each age group. However, if that were to be the case, we should observe a substantial shift in the lifecycle profiles during the great recession (2008), in a similar fashion to the movement observed in Figure 5. Instead, our evidence points to slow moving and more systematic shifts in lifecycle profiles, indicating the more plausible interpretation is that the relationship between consumption levels across age groups has been structurally changing in recent decades.

We conclude that the hump-shaped lifecycle patterns as reported by Attanasio et al. (1999) and Aguiar and Hurst (2013) among others may be a product of pooling that considers lifecycle profiles as being time invariant. We show lifecycle consumption profiles have systematically shifted over the years and thus pooling the data across all households in sample masks changes in the lifecycle behaviour of age groups over time.

2.2.1 Cohort Effects

We now augment the benchmark lifecycle consumption model to include cohorts effects and, as in Aguiar and Hurst (2013), verify that cohort effects do not significantly affect the age profiles depicted above. Let $D_{i,t}^{Norm\_Time}$ be a vector of normalised year dummies following Deaton and Paxson (1994) (restricting year effects to add to zero over the sample period and be orthogonal to a time trend), D_i^{Cohort} be a vector of 1-year birth cohort dummies, and $D_i^{Norm\_Cohort}$ be a vector of normalised birth cohort dummies that similarly add to zero and are orthogonal to a time trend.

For both the *Pooled Lifecycle* model and the *Time-varying Lifecycle* model we estimate both the *Cohort view*, where time effects are orthogonal to a time trend and simple cohort effects are added together with age effects and the *Period view*, where the cohort effects added are normalised, assuming they are orthogonal to a time trend (see Appendix B.5 for details). Formally,

$$c_{i,t} = \alpha_i + \beta_g D_{i,g}^{Age} + \delta_t D_{i,t}^{Norm\_Time} + \zeta_C D_i^{Cohort} + \psi_Z Z_{i,t} + \epsilon_{i,t} \quad \textit{Pooled Cohort View} \tag{5}$$

$$c_{i,t} = \alpha_i + \beta_g D_{i,g}^{Age} + \delta_t D_{i,t}^{Time} + \zeta_C D_i^{Norm\_Cohort} + \psi_Z Z_{i,t} + \epsilon_{i,t} \quad \textit{Pooled Period View} \tag{6}$$

and,

$$c_{i,t} = \alpha_i + \gamma_{g,t}^1 D_{i,g}^{Age} + \gamma_{g,t}^2 D_{i,g}^{Age} \times D_{i,t}^{Time} + \delta_t D_{i,t}^{Norm\_Time} + \zeta_C D_i^{Cohort} + \psi_Z Z_{i,t} + \epsilon_{i,t} \tag{7}$$

Time-varying Cohort View

$$c_{i,t} = \alpha_i + \gamma_{g,t}^1 D_{i,g}^{Age} + \gamma_{g,t}^2 D_{i,g}^{Age} \times D_{i,t}^{Time} + \delta_t D_{i,t}^{Time} + \zeta_C D_i^{Norm\_Cohort} + \psi_Z Z_{i,t} + \epsilon_{i,t} \tag{8}$$

Time-varying Period View

Adding cohort effects to the benchmark *Pooled Lifecycle* model (2) do not qualitatively alter the lifecycle age effects, which remain hump-shaped (Figure 6), confirming the findings reported by Aguiar and Hurst (2013).

We perform the same comparison between our benchmark *Time-varying Lifecycle* model, (Equation 3) and two models including cohort effects, the *Cohort view* and the *Period view*. Figure 7 displays our results. Again, controlling for cohort effects does not alter the systematic time variation in age effects observed in the last decades. We also find that having time controls and time-varying age effects provide a better fit, according to information criteria, than additionally including cohort dummies under either the *Cohort view* or the *Period view* (see Appendix B.5 for details).

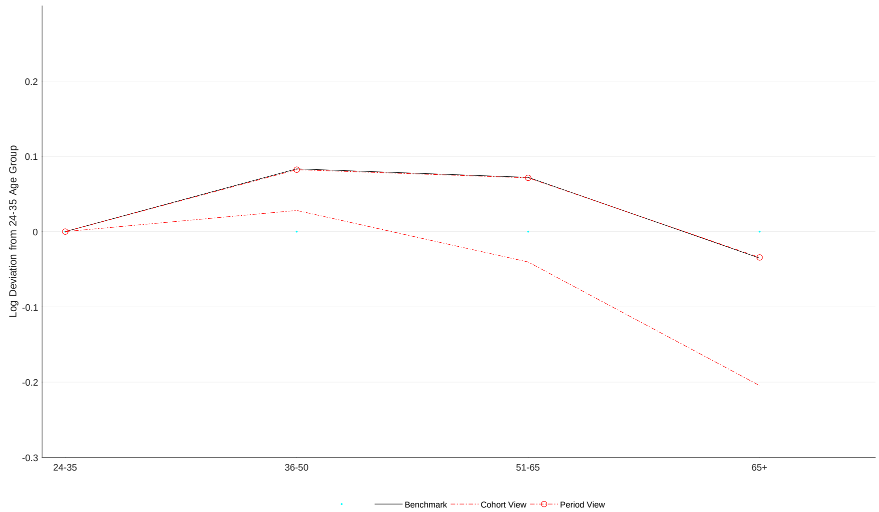

Figure 6: Including Cohort Effects - Pooled Lifecycle

Finally, when including cohort effects we are forced to use orthogonalised time dummies to avoid multicollinearity. Nonetheless, we kept the interaction terms $D_{i,g}^{Age} \times D_{i,t}^{Time}$ with the original time dummies since it allows for a clear comparison between age effects across years maintaining the reference group as the youngest group and the reference year as 1998. We also estimated the model using the interaction defined as $D_{i,g}^{Age} \times D_{i,t}^{Norm\_Time}$. Although under this specification we no longer have a reference year, we obtain similar qualitative results: age effects fall with time and more strongly for the older age groups (see B.5.1 for a more detailed discussion).

Our preferred specification is based on lifecycle profiles viewed from an age perspective rather than cohort. As evident from Figures 1 and 2 we can cut the data both ways. Thus, for the sake of completeness we also estimate our model using age-groups, orthogonalised time dummies and cohort-group effects that can are allowed to vary through time. To do so, we study 5 cohort-groups defined as the household heads with a year of birth between 1918-34, 1935-48, 1949-62, 1963-1976 and 1977-1994.

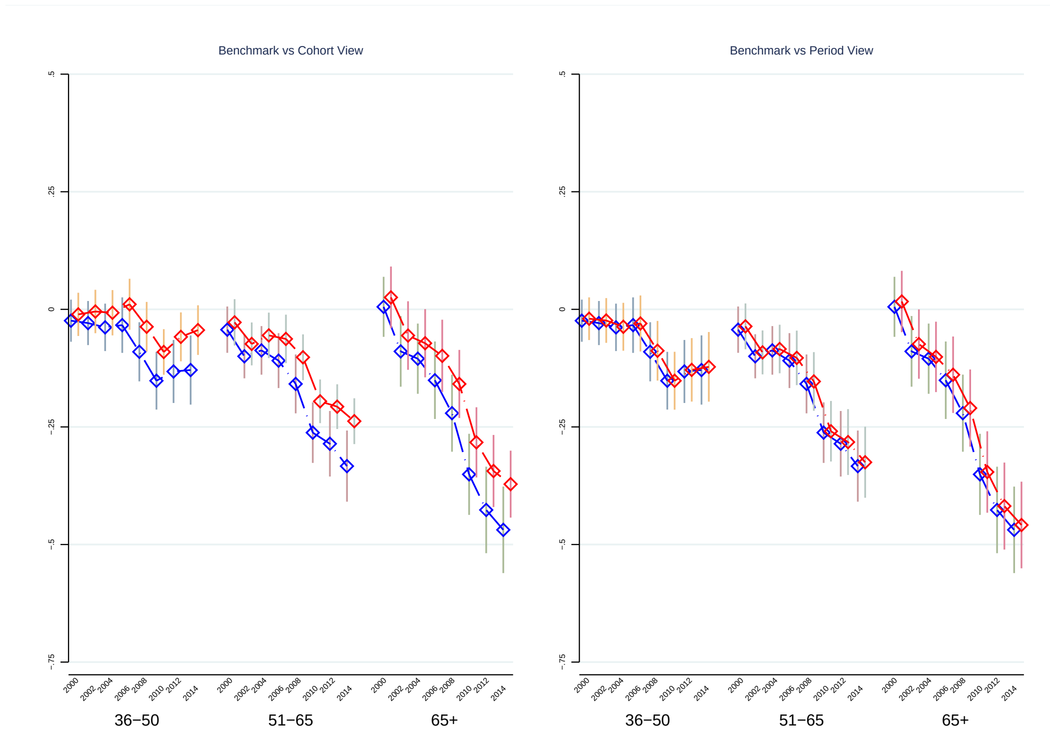

Figure 7: Including Cohort Effects - Time-Varying Lifecycle
$\gamma^2_{g,t}$ - Benchmark (blue), Cohort Models (red)

Formally,

$$c_{i,t} = \alpha_i + \gamma^1_{g,t} D^{Age}_{i,g} + \zeta^1_i D^{Cohort}_{i,g} + \underbrace{\zeta^2_{g,t} D^{Cohort}_{i,g} \times D^{Time}_{i,t} + \delta_t D^{Norm\_Time}_{i,t}}_{Cohort\ Time\text{-}variation} + \psi_Z Z_{i,t} + \epsilon_{i,t} \quad (9)$$

Figure 8 shows cohort effects are also systematically changing through time, with older cohorts facing more variation, confirming that one cannot distinguish an explanation which says consumption in year t is determined by household head's age Age, from an explanation which says it is determined by household head's year of birth, $Cohort = Year - Age$. As such, the key feature emerging from the data is that intergenerational consumption differences have declined from the perspective of age or cohorts. We also observe that the confidence bounds on time varying cohort effects are bigger than the bounds on age effects in the benchmark model. Aside from the more direct interpretation of lifecycle through the lens of age, this confirms the difficulty of estimation the time-varying cohort representation when new cohorts enter and the number of households in the older cohort groups decrease as their mortality rate increases. This problem also prevents us from running the regression with cohort groups with a smaller range of birth years, as we do for age groups of 5 years instead of 15 years.

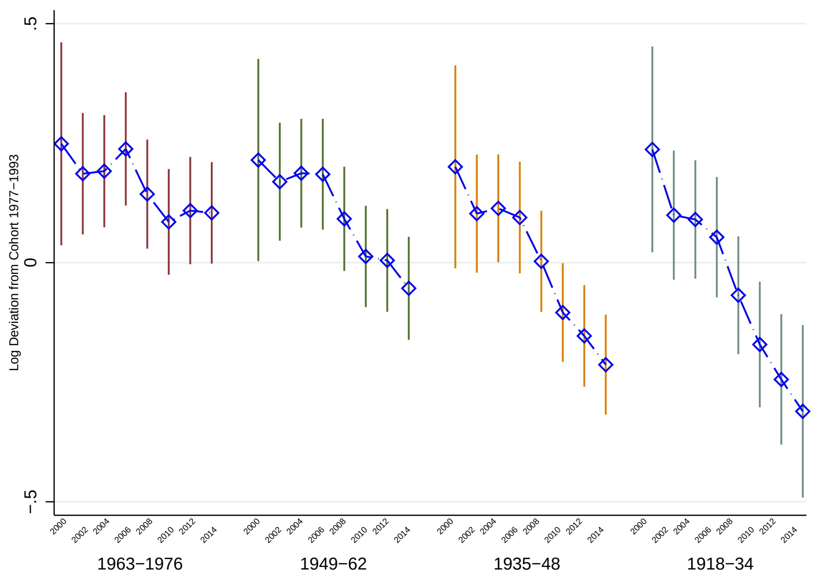

Figure 8: Time-Varying Cohort Effects ($\zeta^2_{g,t}$)

2.2.2 Unobserved Heterogeneity

A crucial aspect of our methodology relates to unobserved household heterogeneity. Our benchmark model leverages the panel dimension of the PSID and estimates lifecycle consumption controlling for household fixed effects (α_i). An alternative approach (see for instance Aguiar and Hurst (2013)) is to estimate the model by OLS when using cross-sectional data such as the CEX. However, ever since the seminal work by Mundlak (1978), it is well known that by estimating with the OLS one cannot control for household unobserved heterogeneity and thus the covariance between age-group (Age_{it}) and α_i may introduce biases in the lifecycle estimates, (β_gs and $\gamma_{g,t}$s). To understand the nature of this bias, we re-estimate *Pooled Lifecycle* and *Time-varying Lifecycle* models with OLS and compare these with the benchmark fixed effect (FE) model. First, we find that for both the *Pooled* and *Time-varying* models, the information criteria strongly favour the fixed effects (FE) approach (reported in the Appendix, Table A.3). Second, inspecting the values of the estimated β_gs and γ_gs over the lifecycle from the FE and OLS estimations reveals significant differences (reported in the Appendix, Figures A.5 and A.7). The OLS lifecycle profiles are very sensitive to the set of household controls included, particularly, house ownership that introduces a level effect and employment status, which affects the shape of the profile at the retirement age-group portion. The fixed effect model produces stable age-group profiles independent of the set of controls, and thus generates robust estimates of the age-group component of consumption. Finally, we find that the OLS *time-varying* estimates no longer exhibit such a clear systematic variation as obtained in the benchmark model.

Aguiar and Hurst (2013) find little difference between OLS and FE for estimating lifecycle consumption profiles using food data in the PSID[12] and conclude unobserved household effects may be safely excluded. We can test the impact of OLS and FE more directly because of the expanded consumption questions introduced in the PSID in 1999. To test whether there is an impact on lifecycle consumption, we first estimate over the food data, following Aguiar and Hurst (2013). Our results confirm that for food data, unobserved household effects have little impact. As mentioned above, when the exercise is repeated with the full nondurable consumption variable, the two

[12]A fuller measure of consumption was not introduced in the PSID until 1999.

approaches are not comparable. We thus conclude that for nondurable consumption, not controlling for fixed effects introduces biases to the age-group parameters and suggests that the assumption that OLS and FE are equivalent cannot be extended to non-food consumption.

2.2.3 Further Robustness

As mentioned in the previous section, we contrasted the OLS and FE models, finding that accounting for unobserved household heterogeneity alters results significantly, while we find that constant cohort effect do not alter our results. Furthermore, we conduct the following robustness exercises.

Comparison with Consumer Expenditure Survey (CEX): The CEX is used in many papers we reference and is generally agreed to provide the gold standard consumption data in the US. However, it is cross sectional (households remain for maximum four quarters) and this rules out controlling for unobserved household level effects. Nonetheless, we reproduce our results using CEX whenever possible. First we re-estimate the pooled model equation 2 with OLS using the CEX and PSID data using comparable observable household characteristics (see the Appendix C for details). Figure 9 plots the resulting age group coefficients (re-scaled to adjust for level effects[13]). The age-group profiles are similar, the correlation between the coefficients is 0.82. The time varying model, equation 4, is also estimated by OLS over both data sets. Although the PSID age-group profiles are a little noisier, they are qualitatively similar; we see the repeated hump shaped lifecycle shapes for both samples. These results indicate that nondurable consumption data in the PSID and CEX have similar life-

[13]The difference is scale is due to the fact that the CEX data are recorded quarterly and the PSID reports annual figures.

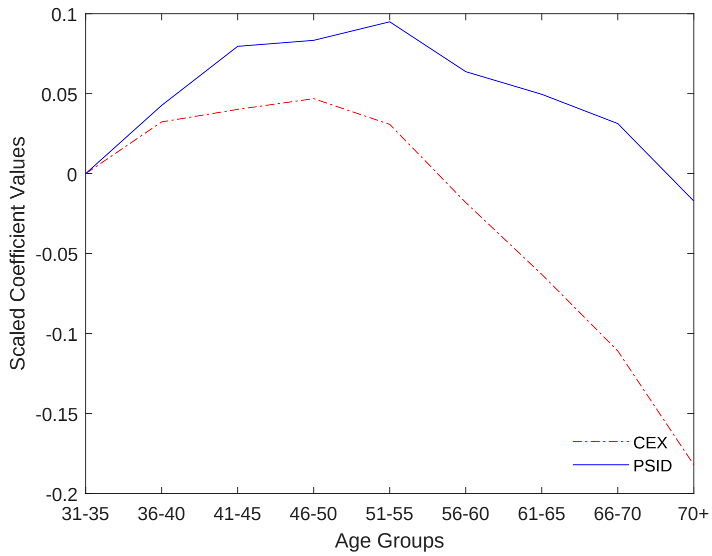

Figure 9: Log consumption over the lifecycle by age group, for the CEX and PSID. Estimated by OLS. The PSID and CEX have been rescaled for comparison.

cycle properties. Thus, we conclude the results presented here are not driven by the potential differences between the PSID and the CEX datasets.

Long Sample: We estimate Equation (4) over a longer sample using an imputed nondurable consumption variable, 1980 - 2014.[14] Whilst the imputation process introduces uncertainty, the results show that the flattening of lifecycle consumption profiles has been occurring since 1980 (see Figure A.11 in the Appendix).[15]

5 Year Age Groups: we re-estimate the benchmark model with 10 age groups, $g = 1, \ldots, 10$ $(24 - 30, ..65 - 70, 71+))$. The systematic changes in consumption lifecycle patterns remain the same, thus averaging the behaviour of households across larger age groups does not alter the main conclusions derived from our empirical evidence. Results are displayed in Figure A.9 in the Appendix.

[14]The imputation method follows Blundell et al. (2008), see the Appendix for details.

[15]As a further check we estimate the model over food data from 1980 - 2014; data on food have been recorded in almost every wave of the PSID since 1968. Time variation in lifecycle profiles are also present. Results from the estimation with consumption of food are available from the authors upon request.

Controlling for Average Income and Subjective House Values: we estimate the model controlling for household's income ($y_{i,t}$) and household's subjective value of housing $H_{i,t}$ (Economic Controls). The modified econometric model is

$$c_{i,t} = \alpha_i + \gamma_{gt} Age_{i,gt} + \gamma_y y_{i,t} + \gamma_H H_{i,t} + \delta_t D_{i,t}^{Time} + \psi_Z Z_{i,t} + \epsilon_{i,t} \tag{10}$$

The time variation in lifecycle consumption patterns is unchanged. Results are displayed in Figure A.12 in the Appendix. We also re-estimate the model using only homeowners. Results remain qualitatively similar.

Consumption Sub-categories: by studying the pooled Consumer Expenditures Survey (CEX) data for the period of 1980-2003, Aguiar and Hurst (2013) show that the lifecycle consumption behaviour for different subcategories are quite distinct. Their findings suggest that possible work-related consumption expenditures, such as clothing, transportation and food away, decline more significantly as households get older. We re-estimate our benchmark model for 9 consumption subcategories in the PSID data (Figure A.13 in the Appendix displays the resulting age group-time coefficient estimates). We observe a flattening of lifecycle consumption profiles in almost all subcategories in our sample, including for the work-related categories, such as transportation and food-away. Thus, our conclusions extend to consumption sub-categories.

Education: the composition of education levels within the population has been changing in the past decades and therefore the time variation we observe could be related to composition effects. To test for this possibility we estimate the lifecycle model for sub-samples of households with different levels of education (i. the grade school only (9.8% of the sample), ii. with high school education (26.9%), iii. some incomplete university education (27.1%) and iv. a university degree or higher (36.1%)).

We find that lifecycle time variation occurs irrespective of the education level. However the observed flattening of lifecycle consumption behaviour is most pronounced by those who have at least high school education. (See Figure A.14 in the Appendix.)

Scaling: we verify the robustness to different ways to adjust for family size, and to deflate consumption expenditures. In the benchmark model we include dummies for number of adults and children but also scaled consumption to reflect family size following Blundell et al. (1994). We estimate the model without scaling, including dummies only, and with scaling but excluding dummies, the main qualitative results are robust to these changes. We test the robustness to different methods of deflation and find that our results are not driven by our choice of using expenditure category specific price indexes. (for details see Appendix B.1).

Family Composition: although results are robust to different scaling methodologies, changes in family composition may be endogenous, potentially introducing selection bias. We re-estimate our model including only stable households (the ones where the head or the spouse did not change). Results once again are qualitatively unaffected. (See Figure A.15)

Panel versus Cross-sectional Estimation: our model makes use of the panel dimension of the data to control for household fixed effects and average (across the sample period) effects of the time varying household characteristics (Z_{it}). An alternative is to estimate the model $c_i = \delta + \beta_g Age_{ig} + \psi_z Z_i + v_i$ for each wave, obtaining a set of β_g's for each wave (t) independently. This model no longer controls for household fixed effects but does allow ψ_z to vary across time. By information criteria the preferred approach for estimation is still fixed effects estimation. Results are shown in Appendix B.6.

2.3 Lifecycle Consumption: The Role of the Income and Housing

We document systematic and significant time variation in the profiles of lifecycle consumption expenditures in the US. Lifecycle consumption profiles have consistently become flatter through time. What may be behind this time variation in the consumption profiles we uncover?

Gourinchas and Parker (2002) stress the importance of the expected growth rate of income in determining consumption behaviour as households age and Attanasio et al. (1999) find that groups of households characterised by a relatively steeper income profile also present a steeper consumption profile, indicating that the shape the income in the lifecycle is a key driver of age group-consumption profiles. Although we have introduced the level of current income into our benchmark model, showing the results are unaffected, relative changes of income across age groups may be relevant in altering the age-group pattern of consumption. Therefore, our first object of interest is the lifecycle variations in income across generations.

Fernandez-Villaverde and Krueger (2011) stress the importance of housing investment in shaping consumption in the lifecycle and indeed for most households investment in a house (typically purchased via mortgage credit) to live-in constitutes the largest asset investment in their lifetime. Moreover, during the first part of our sample, borrowing constraints have relaxed and house prices increased substantially (see Favilukis et al. (2017), Kaplan et al. (2020) and Cox and Ludvigson (2019)).[16]

[16] In the early 2000's there are clear dynamic co-movements between business cycle components of US aggregate consumption expenditures and the Case-Shiller National Home Price Index (which itself is found to be closely linked, at the aggregate level, to our measure of housing wealth, see the Appendix for detail).

Thus, our second variable of interest is the time variation in housing wealth.[17]

In order to extract the role of variations in income and housing wealth across the lifecycle on consumption profiles we allow the age group-time specific components in consumption expenditures not related to household characteristics to vary depending on our variable(s) of interest, namely, household's total family income and subjective housing value. We thus add to our benchmark specification interaction dummies $Age_{i,g,t}X_{i,t}$, that incorporate a variable $X_{i,t} \in \{Y_{i,t}, H_{i,t}\}$ next to our age group-time dummies.

Formally, the econometric model (denoted the *Interaction model*) is

$$c_{i,t} = \alpha_i + \theta_{g,t} Age_{i,g,t} + \theta_{g,X,t} Age_{i,g,t} X_{i,t} + \delta_t D^{Time}_{i,t} + \psi_1 Z_{i,t} + \epsilon_{i,t} \tag{11}$$

To assess the relevance of each of the variable of interest in driving the time variation we can decompose the age group-time effects as follows

$$\gamma_{g,t} = \theta_{g,t} + \theta_{g,X,t}. \tag{12}$$

As such, age group-time dummies ($\theta_{g,t}$) aim to capture age-group specific variation in consumption expenditures that cannot be explained by age-group specific time-variation in our variable of interest (total family income or subjective house value) while $\theta_{g,X,t}$ reflect the contributions of income or housing on the lifecycle consumption profiles.

We report the results in Figures 10 (a) for subjective housing value and 10 (b) for

[17] Although many contributions have looked at the effects of housing wealth in consumption, most have focused on the marginal propensity to consume due to changes in housing wealth (e.g. Carroll et al. (2011), Aladangady (2017), Berger et al. (2018)). In contrast, our interest is in the role of housing in the lifecycle variations of consumption expenditures across generations.

total family income.[18] The top panels depict the age group-time coefficient estimates for the benchmark model and the *Interaction model* by age group, over all time periods. The bottom panels plot the three-way estimates, $\theta_{g,X,t}$, depicting the relevance of housing and income in shaping lifecycle consumption patterns.[19]

First, the role of housing in shaping the lifecycle profile of consumption seems minimal. From the top panel, we still observe the same time-varying lifecycle behaviour of each age group when we control for age group-time specific house valuation; $\gamma_{g,1998} - \gamma_{g,2014}$ and $\theta_{g,1998} - \theta_{g,2014}$ are nearly the same and thus variations in house wealth are not behind the flattening of consumption profiles. The bottom panel shows that $\theta_{g,H,t}$ are generally small, particular for the first 2 age groups. Therefore, high subjective house values seems to sustain consumption particularly for the older households and after the first half of 2000's. Housing wealth seems to be wealth only towards the end of lifecycle and after the 2008 correction (see Buiter (2010)).

In contrast, lifecycle variations in income are more relevant in shaping the changes we observe in consumption profiles. First, from the top panel, in the benchmark model $\gamma_{g,1998} - \gamma_{g,2014}$ increases with age-group, while after controlling for income $\theta_{g,1998} - \theta_{g,2014}$ is fairly constant with age-group. Thus, after extracting the age group-specific component that depend on income, lifecycle consumption are no longer flattening (the only time variation left is a level effect, diametrically opposed to the

[18] In the Appendix (Table A.6) we provide a full description of the estimation results for benchmark and interaction models. i. Time-varying Lifecycle model (Benchmark) ii. Three-way interaction model with Subjective House Value (Interaction SHV model), iii. Three-way interaction model with Total Family Income and finally (Interaction: TFI model) iv. Three-way interaction model with Subjective House Value and Total Family Income jointly (Joint SHV and TFI model). We report coefficient estimates for age group-time dummies as well as estimates for all other controls together with AIC and BIC information criteria.

[19] In both cases the top panels have two y axis to help make visual comparisons between the age group-time coefficients from the benchmark and the Interaction model. We keep the bottom panels with the same y axis to aid in the comparison of the role of housing and income in influencing lifecycle consumption.

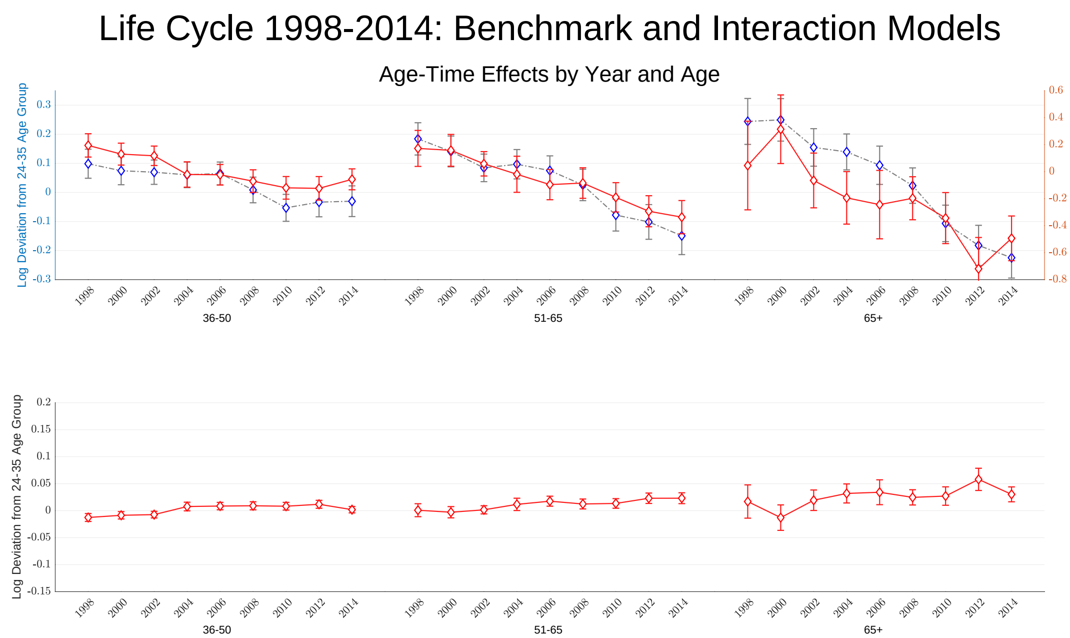

(a) Interaction Model - Subjective House Value

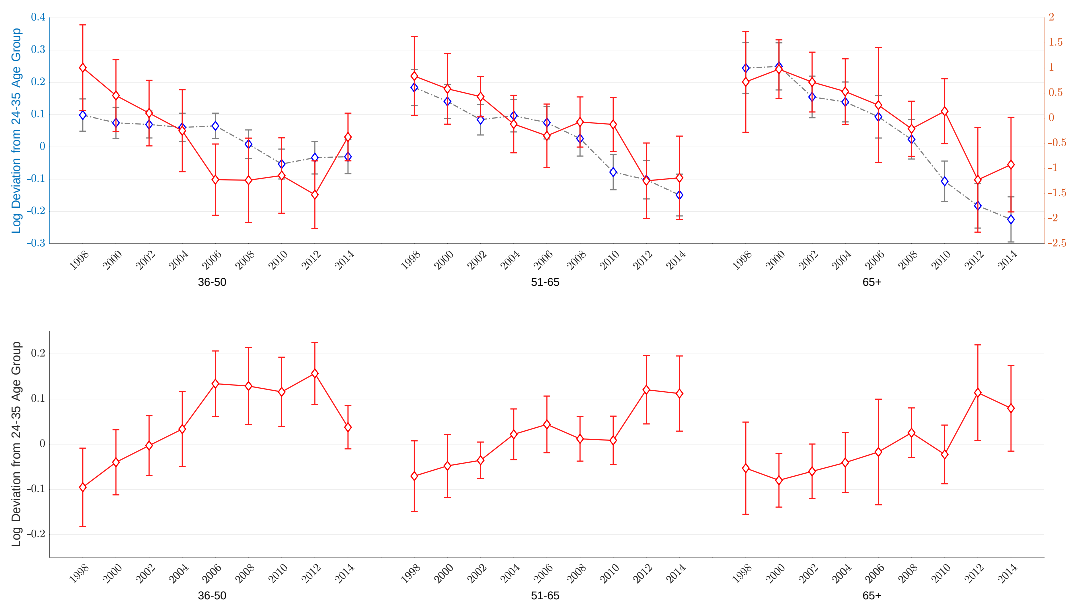

(b) Interaction Model - Total Family Income

Figure 10: Age group coefficients: Benchmark Model (Equation 10) and Interaction Model (Equation 11). **Top** Panel (a) and (b): $\gamma_{g,t}$ from Equation 10 (blue - dash line) and $\theta_{g,t}$ from Equation 11 (red); **Bottom** Panel (a) and (b): $\theta_{g,X,t}$, Equation 11

increasing positive effect of income in driving the age group-profile of consumption). Second, $\theta_{g,H,t}$ increases, indicating that higher income in the lifecycle has become strongly associated with higher consumption levels.

2.4 Time Variation in Lifecycle Income

Our findings so far suggest a close association between time-variation in lifecycle consumption an income. Therefore before we proceed to our theoretical exercise, we complement our lifecycle consumption analysis by presenting detailed patterns in lifecycle income itself. We re-estimate the benchmark model for total family income instead of consumption, extracting the age-group specific path of income for each year ($\gamma^{Y}_{g,t}$).

$$y_{i,t} = \alpha + \gamma^{Y}_{g,t}(D^{Age}_{i,g} \times D^{Time}_{i,t}) + \delta_t D^{Time}_{i,t} + \psi_Z Z_{i,t} + \epsilon_{i,t}. \qquad (13)$$

Our lifecycle income results are displayed in Figures 11 and 12. Indeed we observe a very similar pattern of time variation in income than the one we observe for consumption. After controlling for observable household characteristics, the age group-profile of income has also flattened, with the difference in income across ages decreasing to the point that in 2014 younger households had a higher age group-specific total family income than their older counterparts. We perform the same estimation using labour income instead of total family income and find that the life-cycle income flattening pattern also emerges (See Figure A.17 in the Appendix). Our findings are in line with Kambourov and Manovskii (2009) who report flattening of life-cycle earnings profiles for successive cohorts of male workers entering the labor market in the 1970s and 1980s and related to findings in Jeong et al. (2015) that the lifecycle income (and associated wage premium) flattening may be related to changes in demographic structure. Elsby and Shapiro (2012) also report a decrease in the experience earnings profiles for lower skilled workers. We perform the same robustness exercises as done for consumption and find similar results, the flattening of age-group profiles occur in all specifications (see the Appendix for details). Note again that our results do not imply 35 year old households 2014 are relatively worse off than 35 year old households in the 1998, rather, the results indicate that at each fixed point in time throughout the sample, intergenerational income differences, after fixed effects are accounted for, have decreased in both income and consumption. We conclude the section by stating that the reasons as to why lifecycle income is also flattening is a major research question and requires further structural analysis. We leave to study underlying reasons for further research.

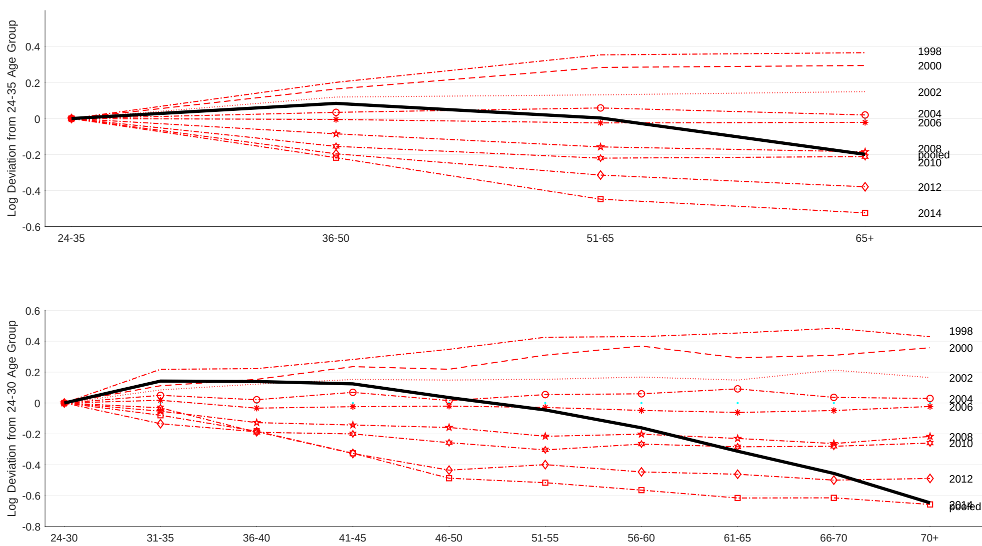

Figure 11: Lifecycle Income Patterns

Note: Each dashed line depicts $\gamma^{Y}_{g,t}$ for each year of the wave of the survey, $(1998, \dots, 2014)$, depicting the estimated lifecycle pattern of consumption for each year. The dark line depicts the age-group effects β^{Y}_{g} when $Age_{i,g,t}$ pools information for the entire sample. The top graph considers 4 age groups, while the bottom graph show the results for 10 age groups.

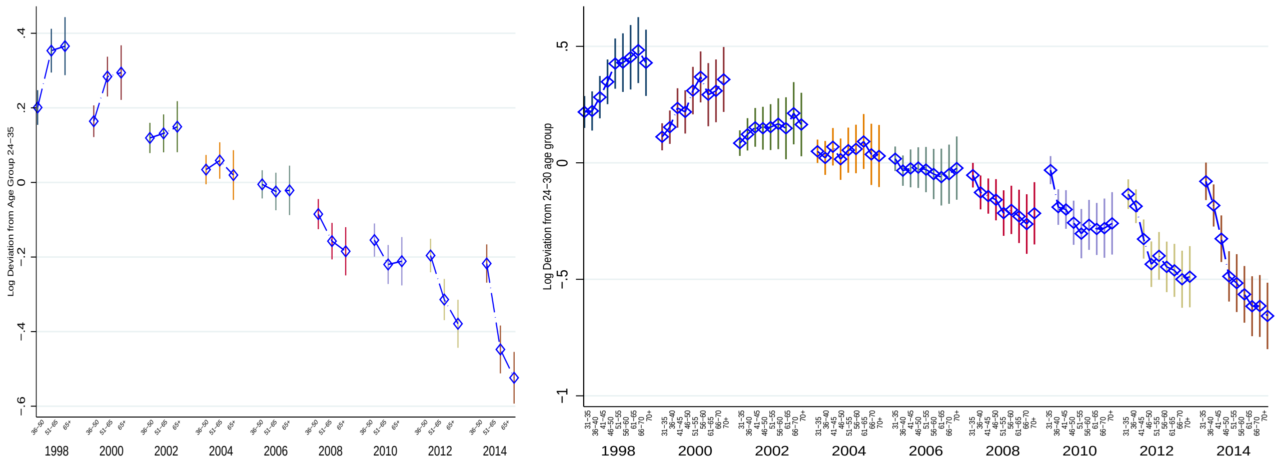

Figure 12: Total Family Income: Age group coefficients ($\gamma^{Y}_{g,t}$) plotted by age group by year. **Left** - 4 age groups, **Right** - 10 age groups.

2.5 Italian Survey of Household Income and Wealth (SHIW)

The results presented so far rely on two modifications to the standard approach for estimating consumption over the lifecycle. First controlling for household fixed effects and, second, allowing consumption to vary by age-group and time rather than pooling data by age-groups over all periods. We show that in the PSID, both of these adjustments have an impact on the lifecycle profile. To investigate whether our results are US specific, we take the model to an alternative panel data set, the Italian Survey of Household Income and Wealth (SHIW). SHIW does not capture every variable we have in the PSID, however the key variables align well with those in the PSID and are sufficient for our purposes. To replicate the structure of the PSID data used in estimation as closely as possible we drop observations before 1998. [20]

[20]Household characteristics are well aligned with the PSID. These are geographical location, retirement status, employment status, educational attainment, marital status, home ownership status. The data are scaled according to household size and composition with the SHIW OECD variable. We also include dummies for number of children and number of adults in the household. Data are deflated as with PSID. Notation for the SHIW variables is consistent with the one used in the rest of the paper. Note also that for the sub sample we use, 28791 households participate in the survey only once. These are dropped from the sample because they will be lost in the demeaning of the fixed effects adjustment. However, the OLS results, where they will not be dropped, are not changed

The SHIW results confirm that the difference between OLS and FE and between constant and time varying age-group profiles are not peculiar to the PSID. Figure 13 depicts the age-group profiles for *pooled* model by OLS and FE, OLS give a stronger *hump* shape over the lifecycle with a peak of the hump in the 50-55 age group. For fixed effects, we observe a flatter profile over the lifecycle, but still displaying the same hump-shape pattern. Results for the *time-varying* FE model for Italian household panel are set out in Figure 14(a) for consumption and Figure 14(b) for income. In both cases we uncover a similar flattening of lifecycle profiles from 1998 to 2012. In 2014 the pattern reverses indicating that for the post-crises the economic conditions of the young in Italy have not been recovering as well as they did in the US (see Glover et al. (2020) for the intergenerational effects of the crisis in the US). In short, the SHIW results confirm that time varying age-group profiles are not peculiar to the US or the PSID.

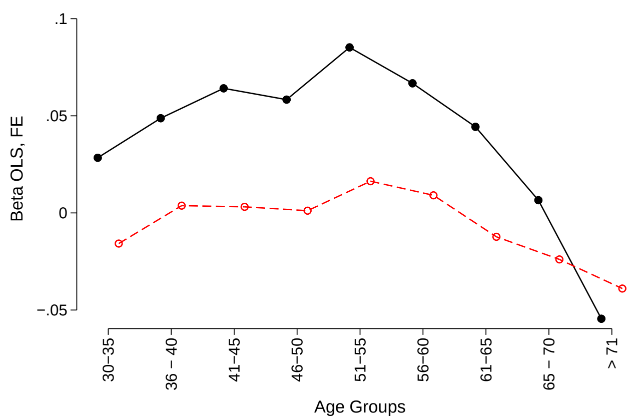

Figure 13: Italian SHIW data: the pooled model. Consumption is the dependent variable. The dotted line plots coefficients from the fixed effects estimation, the solid line, OLS.

---

in any meaningful way by the exclusion of these households. This accounts for the different number of observations in the estimation approach when OLS and FE are compared.

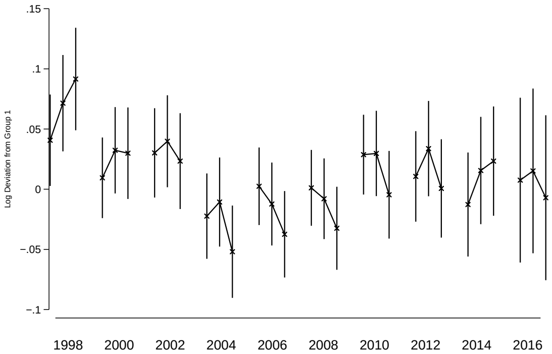

(a) Consumption

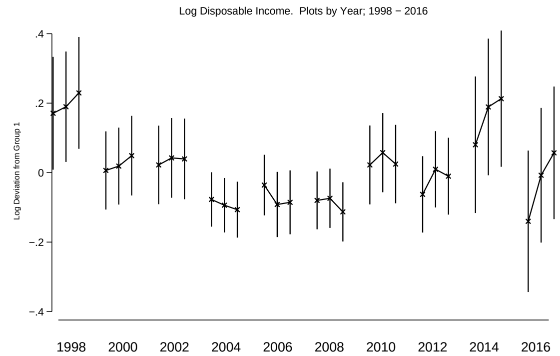

(b) Income

Figure 14: Italian SHIW data: the time varying FE model

3 Lifecycle Model

We now present a theoretical model to gain understanding of the roles housing wealth, credit and income fluctuations may play in driving lifecycle consumption and asset choices. We develop a dynamic, incomplete markets model of household lifecycle consumption and asset choices (following Berger et al. (2018) and Kaplan et al. (2020)). Time is discrete, we set one period of the model to correspond to one year. Population is constant, households enter the economy, work for J_w years, retire and live for another J_r years. A household thus lives for $J = J_w + J_r$ years. Working households face uninsurable idiosyncratic income risk and invest in two assets: a risk-free asset paying a constant interest rate r, and housing. We denote the holdings of each asset by household i at time t as, respectively, $a_{i,t}$ and $h_{i,t}$.

Households born at time t maximize the expected lifecycle utility given by

$$E\left[\sum_{j=1}^{J} \beta U(c_{i,t+j}, s_{i,t+j}) + \beta^{J+1}\mathbb{B}(B_{i,t+J+1})\right]$$

where $c_{i,t+j}$ is nondurable consumption, $s_{i,t+j}$ housing services and $B_{i,t+J+1} = a_{i,t}(1+$

$r) + (1-\delta)P_{i,t+J+1}h_{i,t+J}$ are bequests.

Households are allowed to go short the risk-free asset but must abide by a borrowing constraint. We assume a fraction θ of the current value of owned houses and a fraction ϕ of current income (y_{it}) can be pledged as collateral. Thus, household's asset position must satisfy the borrowing constraint

$$a_{i,t} \geqslant -(\theta P_{i,t}h_{i,t} + \phi y_{it}).$$

Working household's income is given by $y_{it} = exp(\nu(age_{i,t}) + z_{i,t})$, where $\nu(age_{i,t})$ is a known age-group-dependent term and $z_{i,t}$ is a transitory shock that follows an AR1 process. Retirement income is fixed and is assumed to be a function of the income in the last working-age-group period.

Houses are traded at prices $P_{i,t}$.[21] We assume house prices follow a geometric random walk with a drift $P_{i,t} = x_t P_{i,t-1}$, where $ln(x_t) \sim N(\mu_P, \sigma_P)$. μ_P thus denotes the trend growth rate of house prices. Households who trade houses must pay an transaction cost $\Xi P_t h_{i,t}$. Owned houses yield a per-period service equals to $\omega h_{i,t}$, $\omega > 1$, and carry a maintenance cost of $\delta P_{i,t}h_{i,t}$ that fully offsets physical depreciation. Households that decide not to own a house can rent it paying a rental cost of ϕP_t for each unit of housing (the price-rent ratio is constant). Rented houses yield a per-period service equal to $h_{i,t}$, thus owned houses deliver higher services.

At any time t, the household state is fully described by the vector $\mathbf{x} \equiv (a, h, z, P, age)$ given by the liquid asset, housing, income shock, house prices and age. Households face four possible scenarios: (i) household becomes a renter (R), selecting current

[21] Although we include the subscript i, since in our model households may experience different realisations of house prices, these are the prices for the existing house of household i as well as the newly transacted house and in that sense reflect an aggregate shock from the perspective of the household (See Berger et al. (2018)).

housing from the set $\mathfrak{H}^R$ and have no house holdings to carry for the next period; (ii) households that own a house may decide to refinance (F), increasing their borrowing and keeping house holdings $h_{i,t}$ constant, paying a refinancing cost of $\Xi_{Rf}P_{i,t}h_{i,t}$; (iii) household maintains house holdings constant and pays amortization or reduces borrowing (N); and (iv) household is an owner and alters housing stock at time t, or it was a renter in the last period and becomes an house owner (T), selecting housing from the set $\mathfrak{H}$.

Therefore, the value of expected utility of the household is

$$V(\mathbf{x}) = max\{V^R, V^F, V^N, V^T\},$$

where, the value function for each scenario is given by

Renting $V^R(\mathbf{x}){=}\max_{c,a',h'} u(c{,}s){+}\mathbb{E}_t[\beta(V(\mathbf{x'})\vert z{,}P]$ s.t. $c{+}a'{+}\phi Ph'{\leqslant}y{+}a(1{+}r){+}(1{-}\Xi{-}\delta)Ph$ $a'{\geqslant}\phi y, \quad s{=}h', \quad \mathbf{x'}{=}(a'{,}0{,}z'{,}P'{,}age{+}1)$ $h'{\in}\mathfrak{H}^R$	Trading Houses $V^T(\mathbf{x}){=}\max_{c,a',h'} u(c{,}s){+}\mathbb{E}_t[\beta V(\mathbf{x'})\vert z{,}P]$ s.t. $c{+}a'{+}Ph'{\leqslant}y{+}a(1{+}r){+}(1{-}\Xi{-}\delta)Ph$ $a'{\geqslant}(\theta Ph{+}\phi y), \quad s{=}\omega h', \quad \mathbf{x'}{=}(a'{,}h'{,}z'{,}P'{,}age{+}1)$ $h'{\in}\mathfrak{H}$
Refinancing $V^F(\mathbf{x}){=}\max_{c,a'} u(c{,}s){+}\mathbb{E}_t[\beta(V(\mathbf{x'})\vert z{,}P]$ s.t. $c{+}a'{+}{\leqslant}y{+}a(1{+}r){+}({-}\delta{-}\Xi_{Rf})Ph$ $a'{\geqslant}(\theta Ph{+}\phi y), \quad s{=}\omega h, \quad \mathbf{x'}{=}(a'{,}h{,}z'{,}P'{,}age{+}1)$	No Housing Adjustment $V^N(\mathbf{x}){=}\max_{c,a'} u(c{,}s){+}\mathbb{E}_t[\beta(V(\mathbf{x'})\vert z{,}P]$ s.t. $c{+}a'{+}{\leqslant}y{+}a(1{+}r){+}({-}\delta)Ph$ $s{=}\omega h, \quad \mathbf{x'}{=}(a'{,}h{,}z'{,}P'{,}age{+}1)$ $a'{>}$[amort a if $a{<}0$, $0{=}$ if $a{>}0$]

Parameterizations

We assume the per period utility and bequest functions are given by

$$u(c,s) = \frac{1}{1-\sigma}(c^{(1-\alpha)}s^{\alpha})^{(1-\sigma)}, \quad \mathbb{B}(B) = \frac{\psi}{1-\sigma}(B-\bar{B})^{(1-\sigma)}.$$

Households enter the economy with 25 year of age, work for 35 years (J_w), retire and live an additional 20 years (J_r). We set the coefficient of relative risk aversion σ to 2 and the interest rate to 2.4%.

We calibrate the house price process by setting $\mu_P = 0.012$ and $\sigma_P = 0.046$ to match the annual standard deviation and real growth rate of aggregate house prices in FHFA data from 1990 until 2019. We choose a depreciation rate of housing $\delta = 2.2$ to match the depreciation rate in BEA data from 1960 to 2014. The collateral constraints parameter θ determines the minimum mortgage down payment, and we choose a value of 0.8 in our baseline calibration. The ratio of non-collateral debt and income in Survey of Consumer Finance (SCF) in 1998 is around 25% we thus set $\phi = 0.25$. We set Ξ=0.05. This transaction cost is equal to the value of housing adjustment costs calibrated in Díaz and Luengo-Prado (2010).

The working age income process has an age-dependent and a transitory component. Following Floden and Linde (2001), the temporary component z follows an AR1 process with autocorrelation $\rho_z = 0.91$ and standard deviation $\sigma_z = 0.21$ to match PSID earnings statistics (after removing age-dependent components). We calibrate the model using the age dependent component of income estimated for 1998, denoted $\nu(age_{98})$ and depicted in Figure 12. Finally, households receive a social security payment of forty percent of their labour income prior to retirement.

Rented and owned housing are selected within the sets $\mathfrak{H}^R = [0,\ HRmax]$ and $\mathfrak{H} = [Hmin,\ Hmax]$, respectively, where $HRmax < Hmax$. Thus, owned houses cannot be too small and rented houses are in general smaller than owned houses. $Hmax$ is set such that households are not constrained in choosing big houses.

Parameters, $HRmax$, $Hmin$, and α, which controls the share of housing in the utility, β, the discount factor, ψ and $\bar{B}$, which control the bequests, ω, which controls the added utility of house ownership, and ϕ, which controls the rental price, are calibrated to match the ratio of the average earnings of owners to renters of 2.1 (1998 SCF) and the lifecycle profiles of housing wealth, non-housing wealth and homeownership in the 1998 SCF data. We compute average housing wealth and average liquid wealth net of debt for households in nine age groups (25-29, 30-34,..., 60-64, 65 and over). Housing wealth comprises primary residence and other residential and nonresidential real state. Liquid wealth net of debt is the sum of cash, money market, checking, savings and call accounts and holdings of mutual funds, stocks and bonds net of credit cards and mortgages (we only have one asset in the model).[22] For retired households (above the age of 60 years) we also include retirement accounts. In the model, payments from retirement accounts take the form of a lump sum transfer at retirement and a pension annuity, which within our calibration procedure are set, respectively, as fractions ϑ_0 and ϑ_{pa} of the labour income prior to retirement.

Finally, a household enters the economy at 25 years of age with an amount of housing, liquid assets and income such that we match the distribution of 20-30 year old age-group households in the 1998 SCF. Based on our calibration procedure, α, β, ψ, $\bar{B}$, ω, ϕ, ϑ_0, ϑ_{pa}, $Hmin$ and $HRmax$ are:

As in Berger et al. (2018), the model does a good job in matching the SCF asset holdings data.[23] The lifecycle profiles of housing wealth, non-housing wealth and

[22]To normalise data and model we divide both measures of wealth by average income of working age-group households.

[23]In order to solve the model we select a grid of 50 points for assets and housing. To incorpor-

α	β	ψ	$\bar{B}$	ω	ϕ	ϑ_0	ϑ_{pa}	$Hmin$	$HRmax$
0.165	0.9375	2	1.4	1.18	0.05	1.2	0.35	0.1	0.75

Table 2: Parameter Values

homeownership in the data and model are shown in Figure A.23 in Appendix H.

Time Variation in Lifecycle Consumption Profiles

Our empirical results suggest the key driver of the flattening of lifecycle consumption profiles is the change in the lifecycle income profiles. Our benchmark model incorporates the age-group dependent component of income estimated for 1998, denoted $\nu(age_{98})$. As our empirical results show the age-group dependent component of income has been consistently changing from 1998 to 2014, with the difference of total income across age groups decreasing as time passes. We obtain from our estimation two additional age-group dependent component curves, one for 2006, $\nu(age_{06})$, and one for 2014, $\nu(age_{14})$ and re-simulate the model using different age-group dependent income profiles.

First, although not part of our calibration, the model does a good job in matching the age-group profile of consumption observed in 1998. Second, by only changing the age-group component of income $\nu(age)$ we can assess the role of changing income profiles on lifecycle consumption in our model economy. Results are displayed in Figure 15. In all cases we depict the invariant lifecycle profiles for which the age-group income profiles differ but average income remains constant. The theoretical results confirm the empirical evidence that changes in income profiles are crucial to explain the decrease in the difference of consumption across households of different

ate trend in house prices we solve the model such that household select housing wealth $P_{i,t}h_{i,t}$, discounting the continuation value in the Bellman equation by the expected trend in house prices (see Berger et al. (2018) for further details). Invariant lifecycle measures are calculated after we simulated lifecycle decisions for 10000 households.

ages observed in the data. The model is able to match the estimated flattening in consumption profiles reasonably well.

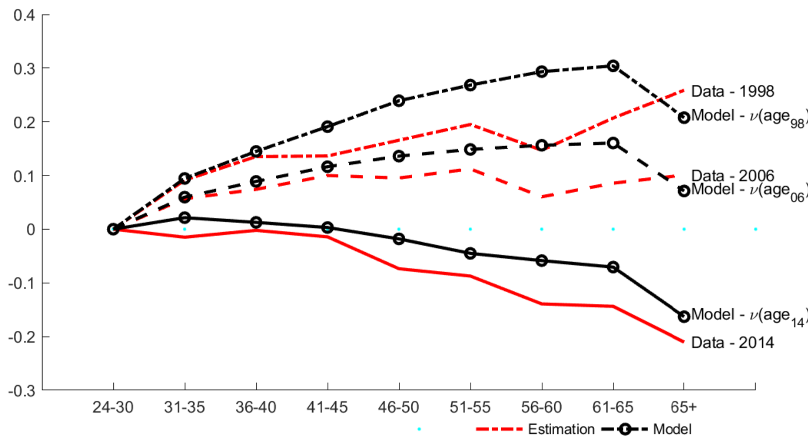

Figure 15: Consumption Life Cycle Profiles: Model with different income profiles versus Estimation

Note: For each model simulation we use either the income profiles $\nu(age_{98})$ - Calibration, $\nu(age_{06})$ or $\nu(age_{14})$ obtained from the estimation (13) - (lines with circles). Data comes from the benchmark estimation of age-group profiles of consumption ($\beta_{g,t}$).

Next we focus on the role of interest rates, credit and housing market dynamics on consumption profiles. From 1998 till 2006 debt to income has increased by 40% (Using data from SCF 1998 and 2007). As we mentioned above, several contributions highlight the importance of relaxed credit constraints during this period. House prices (FHFA data) from 1990 until 2006 increase on average 2.3% as opposed the 1.2%, our calibrated figure, which relies on data until 2019. Finally, several contributions highlight that in the last decades the equilibrium real rate of interest has consistently fallen (see for instance Aksoy et al. (2019) and Del Negro et al. (2019)). To account for these changes in economic conditions from 1998 till 2006 as potential drivers for the movements in consumption profiles we (i) increase the trend in house prices to $\mu_P = 0.023$, (ii) relax credit constraints (a 10% increase in θ - using the SCF of 1998 and 2007, leverage ratios of new house buyers increase by 10% from 1998 till 2007) and (iii) decrease interest rates by 100 basis points. Results are shown in Figure 16. Relaxed credit constraints and lower interest rates imply households borrow more and bring consumption forward, flattening lifecycle profiles. Higher trends in house prices imply home owners become richer during the lifecycle and consumption profiles become steeper. Overall, consumption profiles are not as significantly affected by the level of interest rates, credit and housing market changes as they do when the age-group component of income changes.

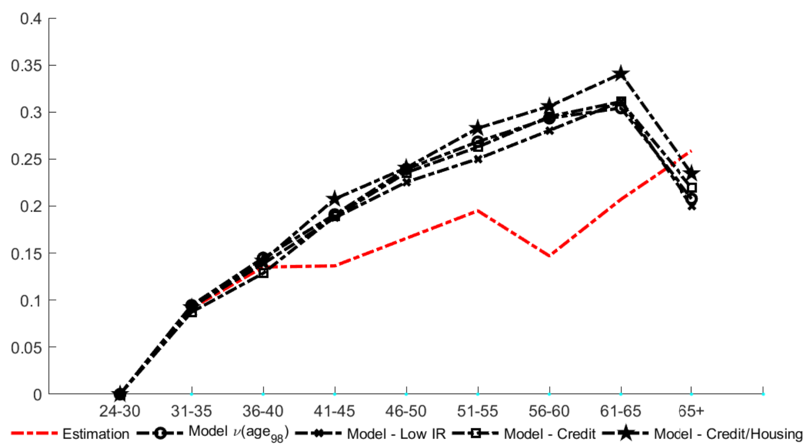

Figure 16: Consumption Life Cycle Profiles: Impact of Credit and House Prices

Note: Model $\nu(age_{98})$ - Benchmark, Model - Credit, incorporates relaxation in credit constraints in the Benchmark Calibration, Model - Credit/Housing, incorporates both the relaxation in credit constraints and the increase in the trend in house prices and Model - Low IR, lowers interest rates in the Benchmark Calibration. Estimation (dash line) comes from the benchmark model of lifecycle consumption for 1998 ($\beta_{g,1998}$)

Time Variation in Assets Holdings and Housing Values

Although the change in the income profiles are sufficient to produce the changes in consumption profiles, we cannot generate the asset accumulation changes observed during the same period. Nonetheless, comparing the lifecycle profiles of housing wealth, non-housing wealth and homeownership in the 2007 SCF and the profiles from the theoretical model incorporating income changes only and income and house prices/credit changes, we show that combining both the changes in income profile and incorporating the changes in the trend in house prices and the relaxation of credit constraints improves the match between data and model both in the changes in consumption profiles and the changes in asset holdings before the Great Recession. Results are reported in Figure A.24 in Appendix H.

In the aftermath of the Great Recession, credit constraints tightened and house prices fell substantially, recovering after 2010. In fact, the average growth rate of house prices from 2007 to 2019 (FHFA data) is close to 0%. To account for these changes in economic conditions since 2007 as potential drivers for the movements in asset holdings we decrease the trend in house prices to $\mu_P = 0$, and tightens credit constraints ($\theta = 0.5$ and $\phi = 0.15$) and compare the asset profiles from the simulated model with the SCF 2013 data (see Figure A.25 in Appendix H). Once again, including only income changes imply simulated asset profiles do not match the data. Incorporating changes in credit and house market conditions help the model in matching asset holdings, although we find the age-group profile of liquid assets under the new income profile portray a much stronger desire to save during the lifecycle as income is no longer expected to increase with age-group. In all cases we depict the invariant lifecycle profiles from the theoretical model. As such, as we compare different simulations all the adjustment/transition process has already occurred. Stock variables such as housing wealth and liquid assets may vary slowly in the data and thus the changes in income profiles we contemplate might take time to affect them. That could be a reason why the model is able to match consumption profiles more closely than the asset lifecycle profiles.

4 Conclusions

We study the evolution of lifecycle consumption patterns using US panel data. We provide extensive empirical evidence that hump-shaped lifecycle profiles of US consumption expenditures are an artefact of pooling data across years from an entire sample and not controlling for household fixed effects. When we account for age-group time interactions not only the hump-shaped profile disappears but we also document clear time varying trends in lifecycle consumption patterns that are robust to a battery of changes in data, in specification and the introduction of additional controls on household characteristics and economic variables. While analysing the potential drivers of this time variation we find that variation in subjective house wealth in the lifecycle do not seem to affect consumption profiles. In contrast, lifecycle income profiles have shown the same time variation and may be behind the systematic variation in consumption we uncover. A lifecycle model of consumption, housing and liquid asset choice shows that indeed changes in lifecycle income profiles are able to generate the observed change in lifecycle consumption patterns. Changes in credit availability and house price dynamics have a much less pronounced effect on consumption in the lifecycle. Nonetheless, in order to also match asset and housing choice, one need to incorporate both changes in income and in housing and credit dynamics. Overall, our results do not imply 35 year old households today are relatively better off than 35 year old households in the 1990s, or that inequality has been changing across time, rather, the results indicate that at each fixed point in time throughout the last decades, the differences across generations have decreased in both income and consumption. Consumption profiles are subject to time variations that can also be interpreted as the result of changes in cohort effects through time.

Our findings complement extensive literature who document widening of overall US and other advanced economies consumption and income inequalities since 1980's (see for instance Aguiar and Bils (2015) and Hoffmann et al. (2020) and references therein). We suggest that observed increase in consumption/income disparities are accompanied by a systematic decline in intergenerational consumption/income disparities and thus may be associated with household specific characteristics rather than age-group or lifecycle effects.

That the flattening of lifecycle consumption in recent decades is associated with flattening of lifecycle income invites further research. We will investigate underlying structural reasons for changing income processes over the lifecycle in our further research.

References

Aguiar, Mark and Erik Hurst, "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 2013, *121* (3), 437–492.

— **and Mark Bils**, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, September 2015, *105* (9), 2725–56.

Aksoy, Yunus, Henrique S. Basso, Ron P. Smith, and Tobias Grasl, "Demographic Structure and Macroeconomic Trends," *American Economic Journal: Macroeconomics*, January 2019, *11* (1), 193–222.

Aladangady, Aditya, "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata," *American Economic Review*, November 2017, *107* (11), 3415–3446.

Attanasio, Orazio and Luigi Pistaferri, "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure," *American Economic Review*, May 2014, *104* (5), 122–26.

— **, James Banks, Costas Meghir, and Guglielmo Weber**, "Humps and Bumps in Lifetime Consumption," *Journal of Business & Economic Statistics*, January 1999, *17* (1), 22–35.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *The Review of Economic Studies*, 2018, *85* (3), 1502–1542.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

— **, Martin Browning, and Costas Meghir**, "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *The Review of Economic Studies*, 1994, *61* (1), 57–80.

Buiter, Willem H, "Housing Wealth Isn't Wealth," *Economics: The Open-Access, Open-Assessment E-Journal*, 2010, *4*, 2010–2022.

Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek, "How Large Are Housing and Financial Wealth Effects? A New Approach," *Journal of Money, Credit and Banking*, February 2011, *43* (1), 55–79.

Cox, Josue and Sydney C. Ludvigson, "Drivers of the great housing boom-bust: Credit conditions, beliefs, or both?," *Real Estate Economics*, 2019.

Deaton, Angus and Christina Paxson, "Intertemporal Choice and Inequality," *Journal of Political Economy*, June 1994, *102* (3), 437–467.

Díaz, A. and M. Luengo-Prado, "The Wealth Distribution with Durable Goods," *International Economic Review*, 2010, *51* (1), 143–170.

Elsby, Michael W. L. and Matthew D. Shapiro, "Why Does Trend Growth Affect Equilibrium Employment? A New Explanation of an Old Puzzle," *American Economic Review*, June 2012, *102* (4), 1378–1413.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 2017, *125* (1), 140–223.

Fernandez-Villaverde, Jesus and Dirk Krueger, "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data," *The Review of Economics and Statistics*, 2007, *89* (3), 552–565.

__ **and** __ , "Consumption And Saving Over The Life Cycle: How Important Are Consumer Durables?," *Macroeconomic Dynamics*, November 2011, *15* (05), 725–770.

Floden, Martin and Jesper Linde, "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, July 2001, *4* (2), 406–437.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose-Victor Rios-Rull, "Intergenerational Redistribution in the Great Recession," *Journal of Political Economy*, October 2020, *128* (10), 3730–3778.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Life Cycle," *Econometrica*, January 2002, *70* (1), 47–89.

Hoffmann, Florian, David S. Lee, and Thomas Lemieux, "Growing Income Inequality in the United States and Other Advanced Economies," *Journal of Economic Perspectives*, November 2020, *34* (4), 52–78.

Jappelli, Tullio and Luigi Pistaferri, *The Economics of Consumption: Theory and Evidence*, Oxford: Oxford University Press, 2017.

Jeong, Hyeok, Yong Kim, and Iourii Manovskii, "The Price of Experience," *American Economic Review*, February 2015, *105* (2), 784–815.

Kambourov, Gueorgui and Iourii Manovskii, "Accounting for the Changing Life-Cycle Profile of Earnings," mimeo, University of Pennsylvania 2009.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Li, Geng, Robert Schoeni, Sheldon Danziger, and Kerwin Charles, "New Expenditure Data in the PSID: Comparisons with the CE," *Monthly Labor Review*, 02 2010, *133*, 29–39.

Mankiw, N.G., "Hall's Consumption Hypothesis and Durable Goods," *Journal of Monetary Economics*, 1982, *10* (Nov), 417–426.

McKenzie, David J., "Disentangling Age, Cohort and Time Effects in the Additive Model," *Oxford Bulletin of Economics and Statistics*, 2006, *68* (4), 473–495.

Mundlak, Yair, "On the Pooling of Time Series and Cross Section Data," *Econometrica*, 1978, *46* (1), 69–85.

Negro, Marco Del, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, "Global trends in interest rates," *Journal of International Economics*, 2019, *118* (C), 248–262.

Schulhofer-Wohl, Sam, "The age-time-cohort problem and the identification of structural parameters in life-cycle models," *Quantitative Economics*, July 2018, *9* (2), 643–658.

St Aubyn, Carolyn, "Measuring the Changing Age Profile of Consumption in the US," Birkbeck, mimeo. 2018.

A Appendix-Data

1. We begin with household heads of the entire Survey, that is 1968 - 2014; there are 270,578 observations. The initial motivation for the PSID was the study of low income households. This original survey is identified as the Survey for Economic Opportunity (SEO). The Survey Research Centre (SRC) later introduced a sample drawn from all income groups and representative of the population. This is the known as the SRC survey and a sample initially of 2,930 households made up this group. In 1990 a new cohort was added to the sample to correctly represent the level of Mexican, Cuban, and Puerto Rican immigrants in the population. Households with income less that zero (64) are dropped. All the variables in the sub sample are truncated at the top and bottom. We convert variables on the truncation boundary to missing. Also heads younger than 25 and older than 80 are dropped. Obvious outliers for food at home, food away from home, food stamps, rent, and from the imputed variable are dropped. The final sub-sample comprises 102,644 observations. There are 11,534 households. The average time in the sample is 8.3 years with a minimum of one and a maximum of 29 years.

2. We consider two measures of housing wealth. Our preferred subjective home value proxy is based on the responses of homeowners to a question in the PSID survey and reported in housing, mortgage distress and wealth data. Ever since the PSID began home-owners are asked what value they attach to their home. Specifically homeowners are asked:

   ‘A20. Could you tell me what the present value of (your/their) (apartment/mobile home/house) is (including the value of the lot if (you/they) own the lot)–I mean about how much would it bring if (you/they) sold it today?'

The question offers an insight into subjective expectations of households about their perceived wealth over a 50 year time period. Household responses to this question define our subjective variable $H_{i,t}$=Subjective Current Home Value. How well do household's subjective home values match prices in the market? In Figure A.1 we compare average values in our sample to the Case-Shiller House Price Index.[24] The two series have a correlation coefficient in the order of 0.96. The relationship holds across house values by income groups; house values in the 10th, 25th, 75th and 95th percentiles have a similar correlation value to the overall value. We plot in Figure A.2 business cycle components of US ag-

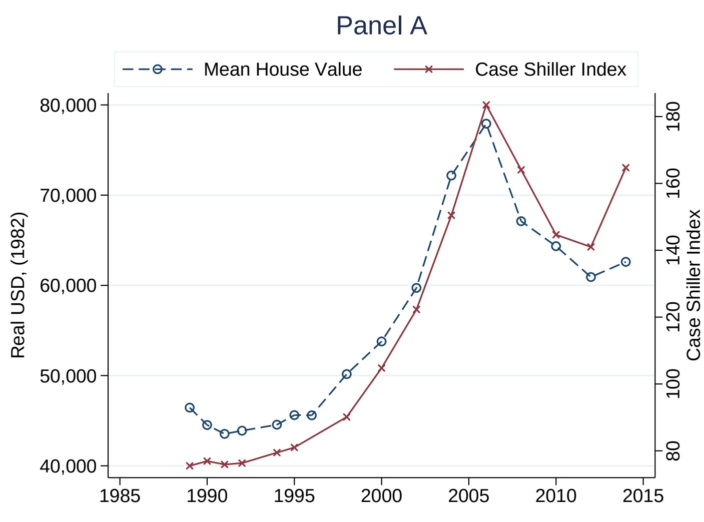

Figure A.1: **S**ubjective House Value and Case Shiller Index

gregate consumption expenditures and Case-Shiller National Home Price Index together. Causal observation suggests that as of early 2000's there are clear dynamic co-movements between these two variables with episodes before and after

[24]This is compiled from repeat sales values of houses in the US across nine census divisions.

the Great Recession particularly marked.[25] The second proxy is the Subjective

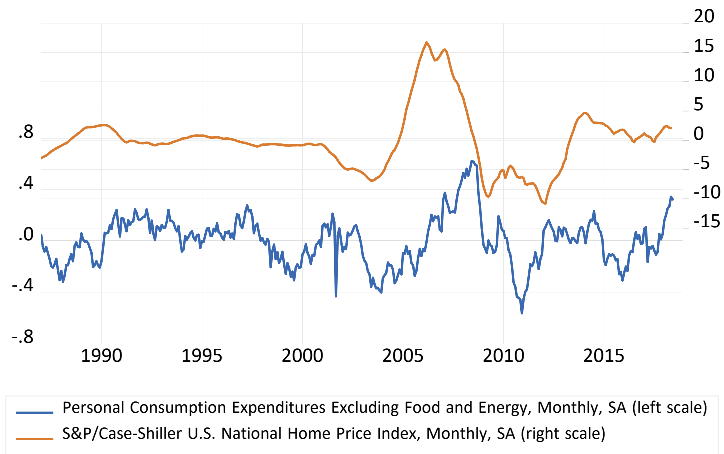

Source: FRED

Figure A.2: Business Cycle Components of Consumption and Case-Shiller House Price Index

Net Home Equity ($HE_{i,t}$) defined as the difference between the subjective house value $H_{i,t}$ and the outstanding mortgage debt ($MD_{i,t}$).

[25]The simple dynamic correlation between 12 month lagged Case-Shiller index and consumption expenditures is in the order of 55%.

Appendix A.1 Unconditional Lifecycle Consumption by Year

We plot the mean consumption by age, by year, $\bar{c}_{age} = \frac{1}{N_{age}} \sum_{age=20}^{80} c_{it}$. There are no controls for household size, composition or any other household level effects. Results are shown in Figure A.3. The typical hump shape over the lifecycle is evident in each of the years.

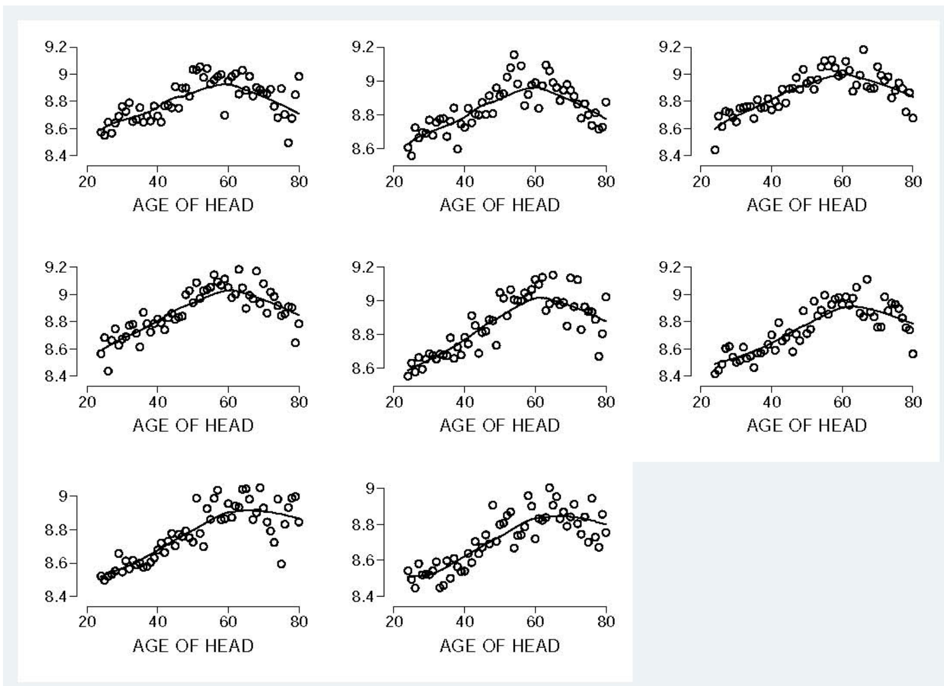

Figure A.3: The panel show mean consumption plotted by age for each year. There are no controls on the data here.

Appendix A.2 Imputation

We again use the PSID, but now include data from 1980 - 1998. The expanded consumption questions were introduced in the 1999 wave. Before this the PSID consistently collected information on a few consumption items: food, home rent and utility payments. For the 1980 - 2015 analysis, we construct an imputed measure of nondurable consumption expenditures following Attanasio and Pistaferri (2014). The imputation approach is based on predicting non-food consumption using an approximate demand system that relates consistently available consumption data (food) to nondurable expenditures.[26] A final adjustment is to drop all odd years to match the biennial structure of the survey after 1998. The final sample consists of 71,662 observations with an average household participation length of 5 waves (or 10 years as we retain biennial waves), a maximum of 17 (16.68 percent), and a minimum of 1 period (2.58 percent). In the long sample homeowners make up about 67.18% of the households.

To estimate imputed nondurable consumption in the pre 1999 data we estimate a log/levels equation by OLS. Specifically, to estimate imputed nondurable consumption in the pre 1998 data we estimate a log/levels equation by OLS in the short sample.

$$\mathbf{nf}_{it} = \mathbf{Z}'_{it}\beta_k + g(f_{it};\theta) + \mathbf{P}'_t\gamma$$

Where

- $nf_{i,t} = \ln(\sum_k C_{it,k})$ is total nondurable, non-food expenditures, with $C_{it,k}$ the expenditure on non-food category k by household i in time t.

[26] Any prediction using this proxy for nondurable consumption expenditures makes assumptions about the stability of relationships between household characteristics and expenditures that we unfortunately cannot test. To limit uncertainty, we choose 1980 as our earliest data point.

- $\mathbf{Z_{it}}$ is a vector of socio-economic variables in the food demand equation.

- g is a polynomial function for f, the total of food at home, away, and the monetary value of food stamps received. These data are available for all waves except 1981 and 1982.

- $\mathbf{P}$ is a vector of annual price indexes; for overall CPI, food at home and food away from home and rent.

Imputed log total nondurable consumption for 1980 - 2014, $\mathbf{\hat{c}_{i,t}}$ is then

$$\mathbf{\hat{c}_{it}} = \log[\mathbf{food}_c + exp(\mathbf{Z}'\hat{\beta} + g(f_{it}; \hat{\theta}_c) + \mathbf{P}'\hat{\gamma})]$$

B Appendix - Specification Issues

Appendix B.1 Scaling of the Data

We investigate how best to adjust for household size and composition. Our findings lead us to control for the number of adults and number of children with dummies and also to use OECD equivalence scales ((Blundell et al., 1994)). We show our results are robust to using only dummies to correct for family size as in Aguiar and Hurst (2013).

As is well documented, family composition influences consumption. Failing to control for family composition distorts the intertemporal pattern of consumption and over states the relationship between consumption and income ((Blundell et al., 1994)). Scaling for family composition explains over half of the hump shape seen in the data over the lifecycle ((Fernandez-Villaverde and Krueger, 2007)). There are different ways of implementing these controls. One method is to scale consumption using one of the available equivalence scales. The scales attach different weights to adults and children and, in some cases, account for economies of scale also; two adults do necessarily require twice the amount of everything. Each scale has benefits and costs (Fernandez-Villaverde and Krueger (2007))[27]. Another approach is to include dummies for numbers of children and adults, or more elaborate versions of this. (Aguiar and Hurst, 2013) have eleven dummies specifying age groups and gender of children.

One point to note is that although it is clear that controlling for composition is important in measuring the age-group profile of consumption, it does not account

[27]See http://www.oecd.org/eco/growth/OECD-Note-EquivalenceScales.pdf or (Attanasio, Banks, Meghir and Weber, 1999).

for the fact that household composition is endogenous. For example, the arrival of children is not usually a complete surprise, nor is their departure. This information is known somewhat in advance and so probably influences spending and savings decisions before the econometric control appears.

In the absence of an agreed approach, we compare results from estimating equations 2 and 4 with consumption adjusted for household composition in three different ways, set out below. There are six cases to consider. Three for the Pooled Lifecycle model and three for the Time-varying Lifecycle model. We take the information criteria as a measure of best fit.

1. 20 Dummies are included in the model to allow for the number of children and adults, but the consumption variable itself is not treated in any way. Attanasio et al (1995) and Aguiar and Hurst (2013) also allow for age and gender of children.

2. OECD equivalence scales.

   These are many equivalence scales to choose from but OECD scales are used in similar work. To apply this, $C_{i,t}$ is divided by the scale value, $scale_{i,t} = 1 + 0.7(n_{i,t} - 1) + 0.5k_{i,t}$, where n is the number of adults and k the number of children. We estimate equations with log values so

$$c^{sc}_{i,t} = ndc_{i,t} - ln(scale_{i,t})$$

   where $c^{sc}_{i,t}$ = log scaled nondurable consumption$_{i,t}$ and $ndc_{i,t}$ is log nondurable.

3. Consumption is adjusted by OECD scale and a full set of dummies are also included. The motivation for this configuration is that after the log transform-

ation, $c^{sc}_{i,t}$ is not equivalent to its levels counterpart $C^{sc} = \frac{C_{i,t}}{Scale_{i,t}}$ and so further controls are needed to capture household composition effects.

We use information criteria to compare model fit in three cases; with dummies only, with OECD scaling only, and with both. Adjusting consumption by OECD scales *and* including separate dummies for numbers of children and adults in the household provides the best fit. Note that this is in spite of the cost of the introduction of 15 additional parameters. The OECD scaling applies a fixed adjustment to each household but this obviously does not completely describe how household composition changes affect consumption. The dummies are more flexible. We note that the model does not account for differences in returns to scale for different expenditure categories as in (Aguiar and Hurst, 2013) or the endogeneity of family composition.

We can test whether the scaling is correct. Define scale as $S_{i,t} = \sum w_i N_i$, some weight w applied to household size and composition. Then the equation has the form $\ln ndc_{i,t} - ln(scale)_{i,t} = \sum \alpha_i N_i$. Or $\ln ndc_{i,t} = \gamma ln(\sum w_i N_i)_{i,t} + \sum \alpha_i N_i$. The hypothesis that $\gamma = 1$ is not rejected so imposing the scaling on the dependent variable is acceptable. This equation brings out the different way that the number in each category influences log consumption; linearly through the dummies and logarithmically through the scaling. If we plot the coefficient values by year, the effects of the different scaling approaches on the lifecycle consumption estimations are very clear. The less restricted approach of using dummies for number of children and number of adults, captures household specific household composition effects left behind by the more restrictive OECD scaling treatment. (See St Aubyn (2018) for details.)

Appendix B.2 Deflating the Data

We compare two methods for deflating the consumption data. The first uses expenditure category specific price indices in order to account for relative price variations and is applied in the short data set. The second more commonly used method utilises the simple CPI across all expenditure categories. We use CPI to deflate the long data set that is imputed. This is because we do not impute category by category, but by the aggregated nondurable consumption variable. For the short data set, we show that results are robust to either deflation method.

In general, consumption data are deflated for lifecycle analysis by a measure such as overall CPI, or a weighted average of price indices. But some work Aguiar and Hurst (2013) deflates by price indexes specific to spending category. We check the impact of deflation approach by these two methods on lifecycle consumption and find it has only a small affect on the outcome (See St Aubyn (2018) for details.)

Appendix B.3 Household Time-varying Controls

We also consider the contribution of the household characteristics included in Z. The initial choice follows other work in lifecycle consumption, variables which are known to affect consumption are included, see Section 2.2. All the are significant except education, which is dropped due to multicollinearity. Education dummies denote maximum education level achieved. There are four of these, the highest is a college degree or higher. The sample has ages 24 to 80. Education level will only change after the age of 24 either in non standard cases of adult education or, in the sub set of graduates. Otherwise, after the age 24, there will be no change. Estimation is by fixed effects. There is not sufficient time variance in the data to estimate the impact of education level. We thus drop education from the model.

Appendix B.4 Reference Group

In Equations 2 - 4 we want to identify age-group effects in time. We have $g \in [1, G]$ age groups and $t \in [1, T]$ time periods.

$Age_{i,g,t}$ is abbreviated as $A_g T_t$ in this Section for ease of notation. $\beta_{g,t}$ can be interpreted as the log difference in average consumption for each age group-time pair, from the reference group. However, there are different ways to parameterise the age-group/time and time effects and this may affect this interpretation. Because the parameters of interest have two dimensions, age-group and time, for the reference group we can drop the first age group in the first time period or we can drop the first age group for all time periods. We want to be able to interpret the $\beta_{g,t}$ coefficients with reference to their own age group and also in a specific time period, i.e. across time in groups (time series) and as lifecycles for different years (cross-sections). We estimate both specifications described above and compare the results;

Case A

Leave out age group 1, for all t, include age group 2 - G for all time periods; $A_2T_1 - A_GT_T$. Include T-1 time dummies, dropping $t = 1$. There are $NT - 1$ parameters.

Case B

Leave out age group 1 in time period 1 only, A_1T_1 and leave out all time dummies. Again we have $NT - 1$ parameters. Comparing the results we find the following:

In Case A, the coefficients of the $T - 1$ time dummies δ_t^A, are identical to the coefficients $\beta_{1,t}^B$ on the $A_1T_2 - A_1T_T$ dummies in model B (where time dummies are excluded). That is

$$\delta_{t+1}^A = \beta_{1,t+1}^B \tag{A.1}$$

The coefficient of A_2T_1, $\beta_{2,1}^A$ in A is identical to the coefficients on $A_2T_1, \beta_{2,1}^B$ in B.

The coefficient of A_2T_{t+1} in model B is equal to the coefficient of A_2T_{t+1} in model A *plus* the coefficient $\beta_{1,t+1}^B$, which by (A.1) is identical to the time dummy in the corresponding period in model A

$$\beta_{1,t+1+i}^B = \beta_{1,t+1+i}^A + \delta_{t+1+i}^A$$

where $i = (0, 1, ..., T - 1)$

In both cases, the base case is $\beta_{1,1}$ and this acts to scale all the other coefficients.

In Case A, the age-group time coefficients are

$$\beta_{g,t}^A = \beta_{g,t} - \underbrace{\beta_{1,t}^B}_{=\delta_t^A} - \beta_{1,1}^B \tag{A.2}$$

for $g = 2, ..., G$ and $t = 2, ..., T$

The δ_t coefficients capture average time effects from the perspective of the omitted age group. Although the time effects affect all age groups together, they nonetheless are a configuration of year effects and the consumption of the omitted age group; the two cannot be disentangled. In our example age group 1 is omitted. If a different age group was left out, the value of the δ_t's would be different.

In Case B

$$\beta^B_{g,t} = \beta_{g,t} - \beta^B_{1,1} \tag{A.3}$$

for $g = 1, ..., G$ and $t = 2, ..., T$

In Case B, the age-group time coefficient includes the value of δ_t; the average time effect plus the omitted age group's consumption.

In summary, from Equations (A.2) and (A.3)

$$\beta^A_{g,t} = \beta^B_{g,t} - \delta^A_t$$

for $g = 2, ..., G$ and $t = 2, ..., T$

$$\beta^B_{1,t} = \delta^A_t$$

for $g = 1, ..., G$ and $t = 2, ..., T$

$$\beta^A_{2,1} = \beta^B_{2,1}$$

Being clear about the effects of different parameterisations is important for interpretation of the age group-time coefficients. In Cases A and B, the base group is always the first age group in the first time period. This is a constant, subtracted from each age group-time coefficient from group 2 - G. The group one coefficients, for the remaining time periods (t+1) - T are a configuration of average time effects and the consumption level of the youngest age group, δ_t. In Case A, these δ_t's are subtracted from the corresponding age group - time coefficients, Equation (A.3) which can then be interpreted as a cross-sectional lifecycle from the perspective of the youngest age group in each of the time periods. Alternatively, organised by age group over time, the coefficients can be interpreted as a time series of consumption by age group from the perspective of the first time period for that age group. If this were not the case then drawing conclusions about the evolution of the β_{gt}'s would be less clear. Thus Case A is selected, noting that although we cannot separate time effects entirely, we can at least narrow it down to the an age group specific response. For further details on the data and specification issues see St Aubyn (2018).

Appendix B.5 Cohort Effects

When measuring the age-group profile of consumption, controls should be included for cohort effects and business cycle effects. The first recognises that some features of lifetime consumption influences are specific to year of birth. The second, picks up shocks that affect the whole population but in a particular time period.

The difficulty here is that $cohort + age = year$. Deaton and Paxson (1994) devised a method to make the columns of the time dummies sum to zero, thus making them orthogonal to the year effects, t. This is a popular approach and is adopted in much of the literature. [28] We define the orthogonalised dummies, $D_{i,t}^{Norm_Time}$ in

[28] To do this, two columns of the time dummies are dropped (coefficients for the first two years can be recovered) and a set of treated time dummies for $t = 3, .., 8$, is defined, dropping the *year* superscript for simplicity, $d_t^* = D_t + (t-1)D_2 + (t-2)D_1$. D_t are the usual dummies for time that equal 1 when the year is t and 0 otherwise.

the model instead of the standard time dummies D_t. Jappelli and Pistaferri (2017) state age, cohort and time effects cannot all be identified without "imposing non testable assumptions". [29]. This first approach is sometime referred to the *Cohort View*, as any time trend is attributed to cohorts and ages, not to time. Another alternative methodology while including all three effects is to define orthogonalised cohort effects, following a similar procedure done for time effects. This is sometimes called the *Period View* as trends are attribute to time (for further discussion see Schulhofer-Wohl (2018)).

We are interested in estimating the lifecycle profile of consumption for the Pooled Lifecycle model (Equation 2) and the Time-varying Lifecycle model (Equation 4) using different combinations of age, time and cohort controls. The first specification is the benchmark model which includes only time and age effects. Then we consider the *Cohort View* which includes cohort, age and normalised time effects, the *Period View* which includes age, time and normalised cohort effects and finally a model with cohort and age effects only, without time effects (*No Time Effects*). As age effects are our primary objective, they are included in all regressions, and the results of the estimated age parameters are presented in the body of the paper. To establish the best specification, the information criteria are compared. Results are reported in Table A.1 for the Pooled Lifecycle, and Table A.2.

Both the Pooled and Time-varying Lifecycle estimations favour including time effects and excluding cohort effects. That confirms our choice of the more parsimonious benchmark model that drops cohort dummies.

[29] Note that this control only captures the additive effect of macroeconomic shocks of time (Jappelli and Pistaferri (2017)), not those where time effects are not additively separable from age. The assumption then is that time effects are the same for all ages. There are other solutions in the literature. For example McKenzie (2006) suggests a second differencing approach. In this paper, we will begin with cohort, age and follow Deaton and Paxson (1994) with orthogonalised time.

Table A.1: Pooled Lifecycle Model:Comparison of Information Criteria for different controls.

	Benchmark	Cohort View	Period View	No Time Effects
AIC	56654.6	56772.6	56686.3	57016.6
BIC	57367.6	58111.7	58034.1	58294.8
df_m	81	153	154	146

Table A.2: Time-Varying Lifecycle Model:Comparison of Information Criteria for different controls.

	Benchmark	Cohort View	Period View	No Time Effects
AIC	56483.1	56533.5	56514.2	56543.0
BIC	57874.3	58550.9	58540.2	58499.4
df_m	159	231	232	224

B.5.1 Time variation with Orthogonalised Time Dummies

When including cohort effects we are forced to use orthogonalised time dummies to avoid multicollinearity, while keeping the interaction terms $D_{i,g}^{Age} \times D_{i,t}^{Time}$ with the original time dummies. That allows for a clear comparison between age effects across years maintaining the reference group as the youngest group and the reference year as 1998. We also estimated the model using the interaction defined as $D_{i,g}^{Age} \times D_{i,t}^{Norm\_Time}$ and find similar results: age effects fall with time and more strongly for the older cohorts. In this specification, due to the normalisation, we lose one year of comparison (2000, besides the original base year 1998, since we impose an additional restriction) and the base year becomes a timeless average of all years and thus the value of the age effects obtained are not directly interpretable as before. However, we can still compare the changes of the estimates across the remaining years, from 2002 to 2014 or across $D_{i,t}^{Norm\_Time}$ for $t = 3$ onwards, to draw inference about their relative change. We find similar results, age effects fall with time and more strongly for the older cohorts.

Formally,

$$c_{i,t} = \alpha_i + \gamma^1_{g,t} D^{Age}_{i,g} + \gamma^2_{g,t} D^{Age}_{i,g} \times D^{Norm\_Time}_{i,t} + \delta_t D^{Norm\_Time}_{i,t} + \zeta_C D^{Cohort}_i + \psi_Z Z_{i,t} + \epsilon_{i,t} \quad \text{(A.4)}$$

Time-varying Cohort View - Modified Interaction

Through the orthogonalisation procedure the $D^{Norm\_Time}_{i,t}$ of $t \geqslant 3$ becomes a variable with zeros for $\tau > t$, ones for $\tau = t$ and takes the value of $t-2$ and $1-t$ for year one and two respectively. Thus, $\gamma^2_{g,t}$ leverages information from households in age group g from years t, 2 and 1 in the sample, not allowing for a direct link between its value and year t only. Moreover, the reference age effect is no longer the youngest group at time 1, but the youngest group at an average timeless year. However, as we compare $\gamma^2_{g,t}$ with $\gamma^2_{g,\hat{t}}$ for $t, \hat{t} \geqslant 3$, the relevant change leveraged for identification is the information from group g at year t versus the information from group g at year $\hat{t}$, thus the changes in $\gamma^2_{g,t}$ provide information on whether the age effects are changing for this age group from 2002 to 2014 in the sample. As depicted in Figure A.4 we observe a qualitatively similar time variation in age effects from 2002 onwards as in the benchmark model, confirming that our general conclusions also hold in this specification.

Appendix B.6 Different Approaches to Estimating Consumption over the Lifecycle

Consumption over the lifecycle can be estimated in different ways. What effect do different approaches have on the results? As an additional robustness check, and because the estimation method is important to our results, we consider the effect of alternative specifications.

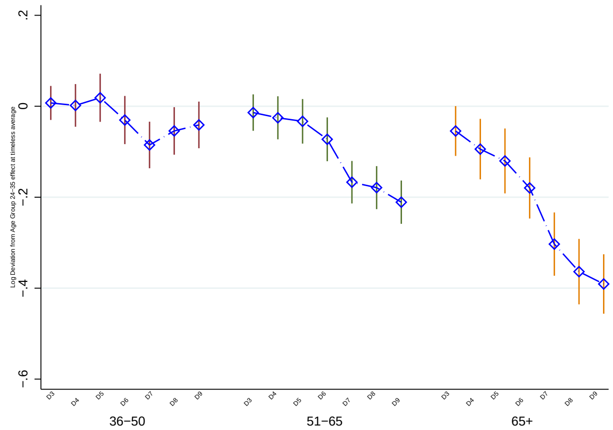

Figure A.4: Time-Varying Lifecycle - Iteration with Orthogonalised Time Dummies

Three models and estimation approaches are considered here. We estimate Equations 2 and 4. Because the data are in panels, we can compare estimation results from 1) pooling all the age-groups over all time periods (Pooled Lifecycle model), and 2) interacting age-group and time (Time-varying Lifecycle model). We can separate this approach in another way, 3) estimating the age-group effects by year as a cross-section (Repeated Cross-sectional model). Pooled Lifecycle models and Cross-sectional models are commonly used in the literature. Both Pooled Lifecycle and Time-varying Lifecycle models can be estimated by OLS or by fixed effects that differ in their treatment of unobserved household effects. We can therefore differentiate the impact of controlling for unobserved household effects, which are likely to be correlated with age-group by inspecting the fixed effects versus OLS estimates.

B.6.1 Pooled Lifecycle Model: OLS versus Fixed Effects

OLS Estimation

This is a standard approach in estimating consumption over the lifecycle. The households in each age group in every time period are pooled and the average effect estimated by β_g.

Estimating by OLS means there are no controls for unobserved household level effects, α_i so the residuals take the form $v_{it} = \alpha_i + \epsilon_{it}$. It is likely that they will be correlated with age-group; $cov(Age_{ig}, \alpha_i) \neq 0$. This means the estimators will likely be biased. If the covariance of Z_{it} with Age_{ig} is not zero this will also effect the value of the β_g's.

Fixed Effects Estimation

Now unobserved household effects can be controlled for. The approach means that the α_i's are subtracted out of the data. This removes any bias in the nine β_g's that resulted from $cov(Age_{ig}, \alpha_i)$. The remaining impact of this estimation approach is a scaling effect on all the variables that change over time. The fixed effects procedure is to subtract the mean effect over all time periods from each observation, $x_{it}^{FE} = x_{it} - \frac{1}{T}\sum_{t=1}^{T} x_{it}$.

Figure A.5 display the results. Both OLS and FE deliver similar age-group profiles of nondurable consumption and food consumption, although for nondurable consumption which includes a more varied set of spending categories that can be much more discretionary, accounting for unobserved household heterogeneity has a more significant impact.

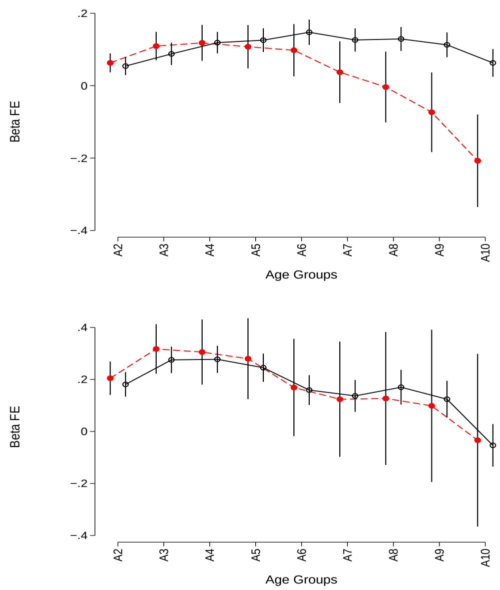

Figure A.5: Estimates of age-group effects for nondurable consumption (above) and food (below). Age groups are pooled over all time periods, 1998 - 2014. Fixed effects (dashed line) and OLS (solid line).

B.6.2 Time-varying Lifecycle: OLS versus Fixed Effects

We are estimating Equation 4, the Time-varying Lifecycle model. Each estimate is the difference with reference to the youngest age group in the first period.

1. By OLS, where we do not control for unobserved household effects.

2. By fixed effects, where we do control for unobserved household effects and where $v_{it} = \epsilon_{it} + \alpha_i$.

Figure A.6 shows coefficient plots from estimation of the pooled model over the PSID with three different specifications of household level controls. We restrict controls to exclude employment status, education, house ownership and state of residence. And then add each of these into the model. We find that the OLS coefficients are very sensitive to these changes. Fixed effects estimation is not, it remains very stable as we change the controls. [30] This exercise reinforces the importance of controlling for household unobserved effects when estimating consumption over the lifecycle.

Two of the controls are particularly relevant for the OLS estimation. House ownership has a sizable level effect on age-group profiles while employment status affects the final part of the lifecycle, due to retirement decisions. Aguiar and Hurst (2013) discuss this finding in detail, linking the fall in consumption at later ages due to the lack of work related expenditures. The FE model, by controlling for all unobserved characteristics, eliminates the biases on the estimated age-group profiles generated due to the lack of household specific controls, delivering robust lifecycle profiles. For food consumption controls do not significantly affect the results. Finally, when we disaggregate the profiles by time, there are systematic cross year differences within the fixed effects estimations that are not obvious in the OLS specification, Figure A.7. The fixed effects pattern pivots over the years from an upward slope to a downward one.

[30] Carrying out the same exercise over the CEX with education shows that the CEX is less sensitive to this. One explanation is its size - the CEX has around 10 times the number of observations that the PSID.

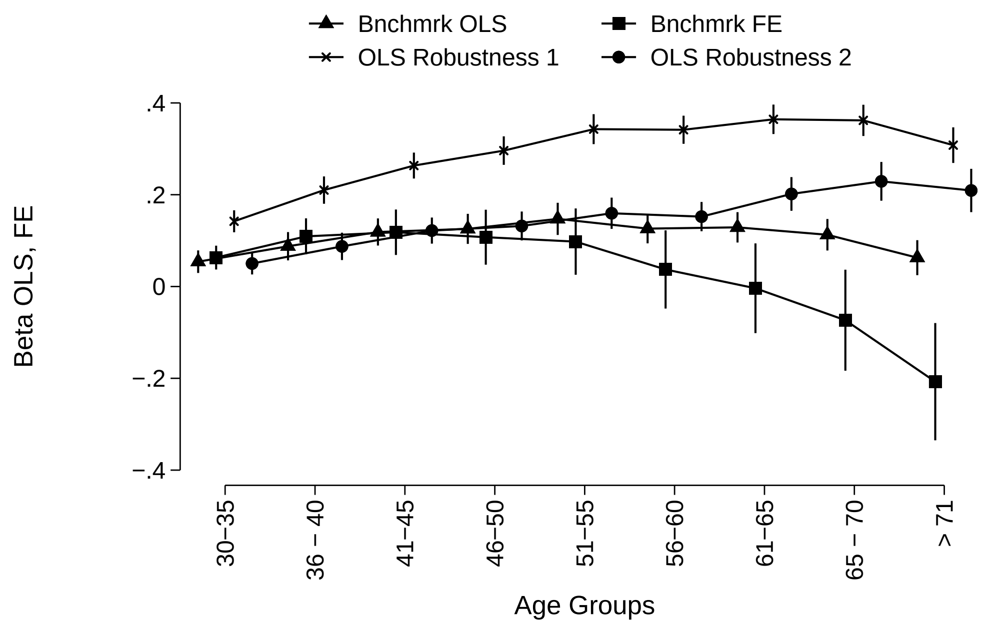

Figure A.6: OLS Robustness with respect to Controls: Both Benchmark FE Model and Benchmark OLS Model include benchmark controls (number of adults, number of children, time, race, education levels, home ownership, self-employment, disability, marital status, and state of residence). In OLS Robustness 1, home ownership is excluded from benchmark controls. In OLS Robustness 2, employment status (inc. retirement) is added to benchmark controls.

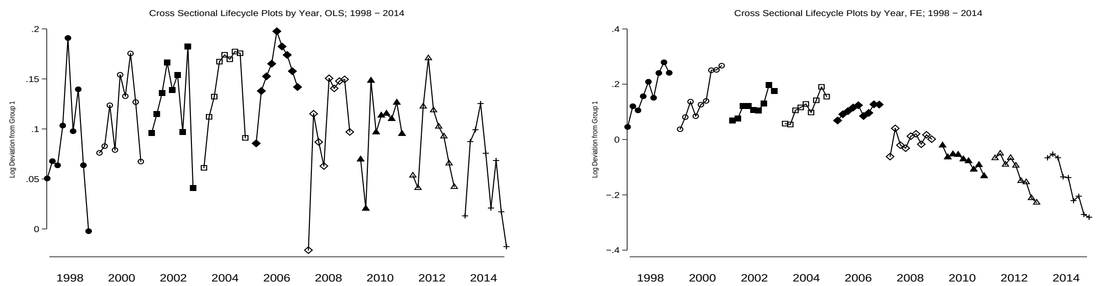

Figure A.7: β_{gt}^{OLS} estimates (left panel) and β_{gt}^{FE} estimates (right panel)

B.6.3 Repeated Cross-sections

We estimate Equation 2, by OLS over repeated cross-sections, i.e. T sets of estimates, one for each time period. Each estimation yields $G-1$ coefficients, $\beta_g t$.

There are three sorts of bias that can arise in estimates of β_{gt} from this approach. The vector of controls in Z_{it} varies in each time period, so rather than estimating their effect as an average over all time periods it is an average for one time period. The covariance $cov(Age_{igt}, Z_{it})$ may or may not equal zero and this may vary in each period and thus in each set of estimation results. Second, we cannot control for unobserved household effects α_i and these are very likely to be correlated with age in each year. The above introduce bias in β_{gt}. Third, we cannot control for average time effects in this approach. If correlated with the Age_{igt}, this will also bias the estimators β_{gt}.

The plots of the estimated coefficients are not reported here but show that although there is variation from year to year, the overall shape of the lifecycle plot is sloping upwards with age-group. This is consistent with the estimation of Pooled Lifecycle model estimated with OLS above. The difference between these two approaches is the covariance of the unobserved household effects, ϵ and controls, Z_{it} with Age_{igt}. Information criteria are reported in Table A.3.

B.6.4 Comparison by Information Criteria

Various models estimated here are sometimes nested versions of each other and sometimes not. One way of comparing all of them, regardless of the structure and relationship, is by information criteria.

For the repeated cross-sections estimation the AIC and BIC are summed for each time period. The sum of the individual information criteria is an appropriate comparison to the AIC and BIC of the benchmark model estimated by the fixed effects specification. Table A.3 displays our results.

Model	obs	LL Null	LL Mod	df	AIC	BIC
Equation 4: FE, Time-varying,	44149	-29602.2	-28081.5	160	56483.07	57874.33
Equation 4: OLS, Time-varying	44149	-52930.1	-45129.5	164	90587.02	92013.06
Equation 2: FE, Pooled	44149	-29602.2	-28174.8	88	56525.66	57290.84
Equation 2: OLS, Pooled	44149	-52930.1	-45167.7	92	90519.4	91319.37
Equation 2: Repeated Cross-sections						
Year						
1998	4156	-3614	-2710	84	5588	6120
2000	4395	-3908	-2888	84	5943	6480
2002	4557	-4239	-3033	84	6234	6774
2004	4605	-4489	-3159	85	6489	7036
2006	4706	-5001	-3711	86	7593	8148
2008	4832	-5141	-3934	86	8040	8597
2010	4838	-4972	-3735	86	7642	8199
2012	4833	-5167	-3933	84	8033	8578
2014	4763	-4925	-3688	83	7541	8078
Total					63104	68010.99

Table A.3: Information criteria for the different approaches for estimating the lifecycle consumption profile.

C Data Description, Consumer Expenditure Survey (CEX)

CEX Non durable consumption variable

As a robustness check, in section 2.2.3 we estimate the pooled model by OLS over data from CEX. We detail the composition of the nondurable consumption variable and household level controls below.

We use quarterly data from the CEX from the time periods to match the PSID data; Q1 1999 - Q4 2014. The expenditure data are scaled for household composition using the OECD scales, as set out in section Appendix B.1. The data are also deflated as set out in section Appendix B.2.

We construct nondurable consumption to match the PSID version as closely as possible. There are some differences which are detailed here.

PSID	CEX
Food	Food[31]
heat, electricity, other	energy, phone bill
water	water
Medical costs, doctor prescriptions, hospital,nursing home	health expenditures
Child care	babysit
bus,cabs, parking	public transport
Vehicle repair, insurance servicing; additional vehicle costs	Vehicle expenditure - services
rent - rent or 6% house value	rent paid, mortgage interest property tax[32]
-	life insurance
gasoline	gasoline
health insurance	health insurance
house insurance	house expenditures (services)

Table A.4: Non durable consumption composition in the PSID and CEX.

Household level controls in the CEX closely match our PSID baseline controls, set out in section 2.1. These are, number of adults, number of children in the household (under 18), a dummy for married households, being a home owner, completed education. We do not include a dummy for state of residence.

Not all our results can be reproduced in the CEX. We are restricted to cases where there is no need for fixed effects estimation. We compare 1) unconditioned consumption over the lifecycle 2) age group coefficient values from estimating our pooled model by OLS 3) coefficient values for the time varying model, estimated by OLS, over the PSID and CEX. In each case considered, both data sets give similar results.

First, comparing deflated and scaled nondurable consumption by age, without additional household controls, confirms the general life cycle shapes are similar for both data sets (results not shown). Because the CEX data are recorded quarterly and the PSID reports annual figures, there is a difference in scale. The CEX shows a lifecycle peak a little earlier and also a steeper fall in consumption after retirement.

We re-estimate the pooled model (Equation 2), over log nondurable consumption in both data sets. Household controls are aligned closely although there are some differences.[33] Figure 9 plots the age group coefficients from these estimations. There are some differences in scale, addressed for comparison by rescaling the coefficients, and standard errors (not shown). Differences in standard errors are to be expected given the relative sizes of the data sets. The correlation between the coefficients is 0.82.

The time varying model, case 2, (Equation 4), is estimated by OLS over both data sets. Although the PSID plots are a little noisier, we see the repeated hump shaped lifecycle shapes in both cases.

[33] The controls for the CEX are set out in C

These results go some way to showing that both the CEX and PSID yield similar results over the lifecycle when estimating with OLS and ignoring household level effects. [34]

[34] Information on the composition of the nondurable consumption and the controls used in the CEX estimations here, are set out in appendix C.

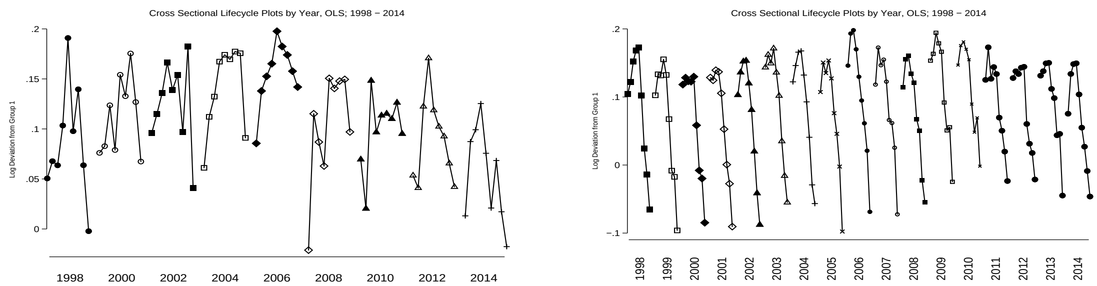

Figure A.8: $\beta_{gt}^{OLS_PSID}$ estimates (left panel) and $\beta_{gt}^{OLS_CEX}$ estimates (right panel)

D Additional Robustness Results

We run several specifications for robustness. Figures not presented in the main text are presented here. We report the estimates of the benchmark model with 10 age groups (Figure A.9) and the results of the estimation of the benchmark model using a sample of homeowners (Figure A.10), using the long sample with imputed consumption data (Figure fig:robust2), introducing controls for household specific economic variables, income ($y_{i,t}$) and household's subject value of housing $H_{i,t}$ (Figure A.12), looking at consumption subcategories (Figure A.13), and partitioning the sample to look at different levels of education (Figure A.14)and to focus on households where the head or the spouse do not change over time (Figure A.15).

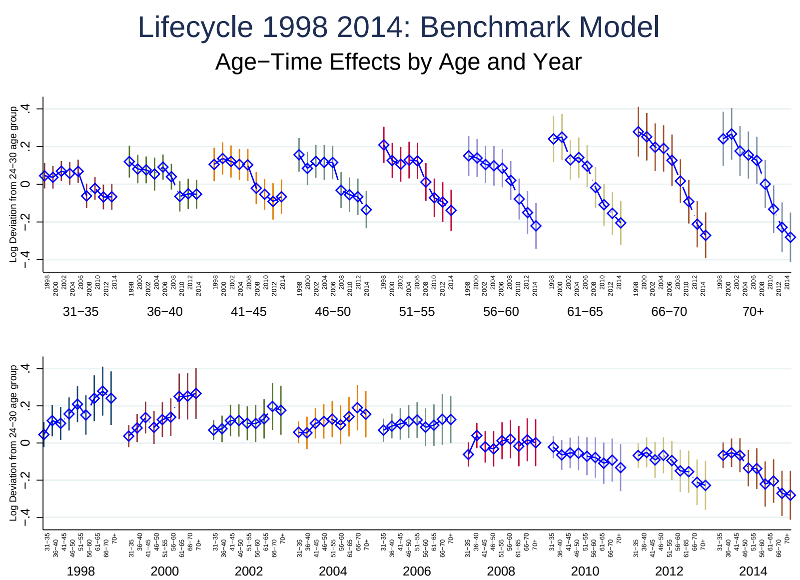

Figure A.9: Model with 10 age groups - Full Sample: **Top** Coefficients by age group. **Bottom** Lifecycle plots by year.

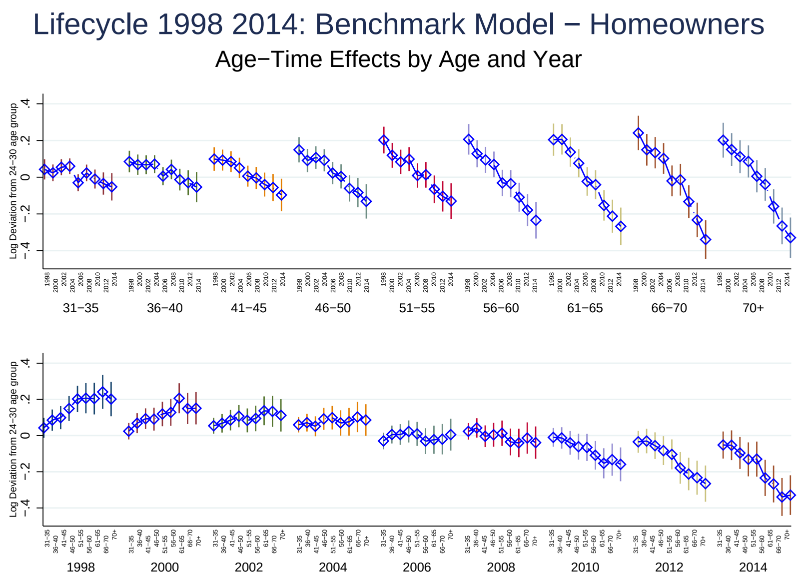

Figure A.10: Model with 10 age groups - Homeowners: **Top** Coefficients by age group. **Bottom** Lifecycle plots by year.

E Additional Results - Interaction Models

We report the results of the interaction model when both house value and income are included.

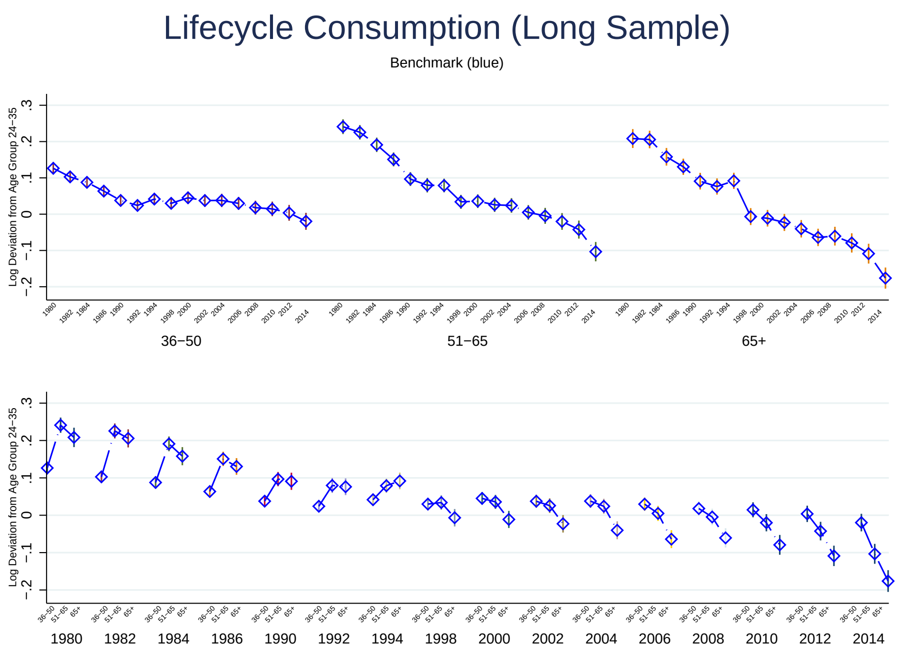

Figure A.11: Age group coefficients: Results from Long Data Set. **Top** Coefficients by age group. **Bottom** Lifecycle plots by year.

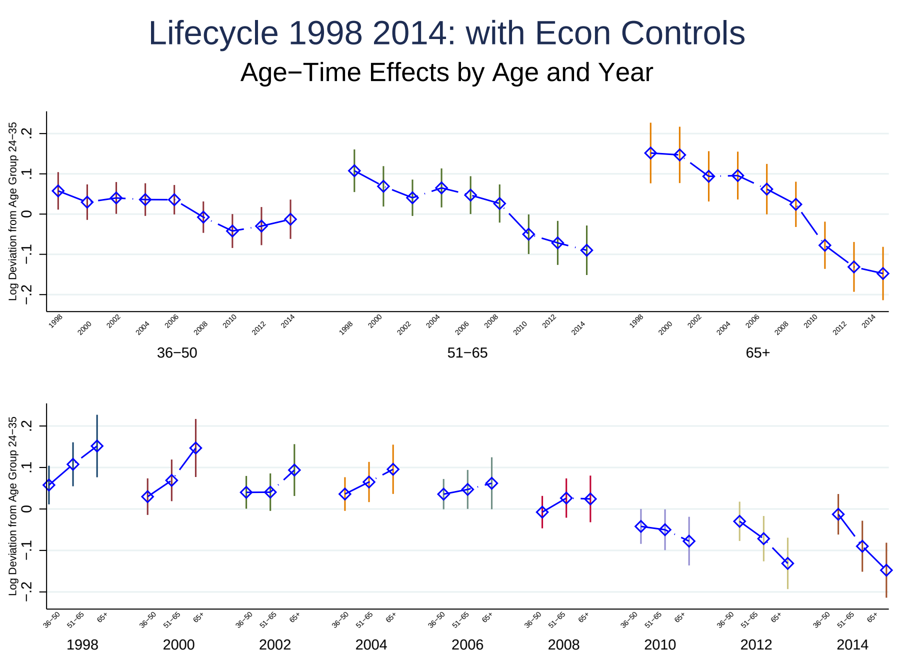

Figure A.12: Model with additional economic controls: **Top** Coefficients by age group. **Bottom** Lifecycle plots by year.

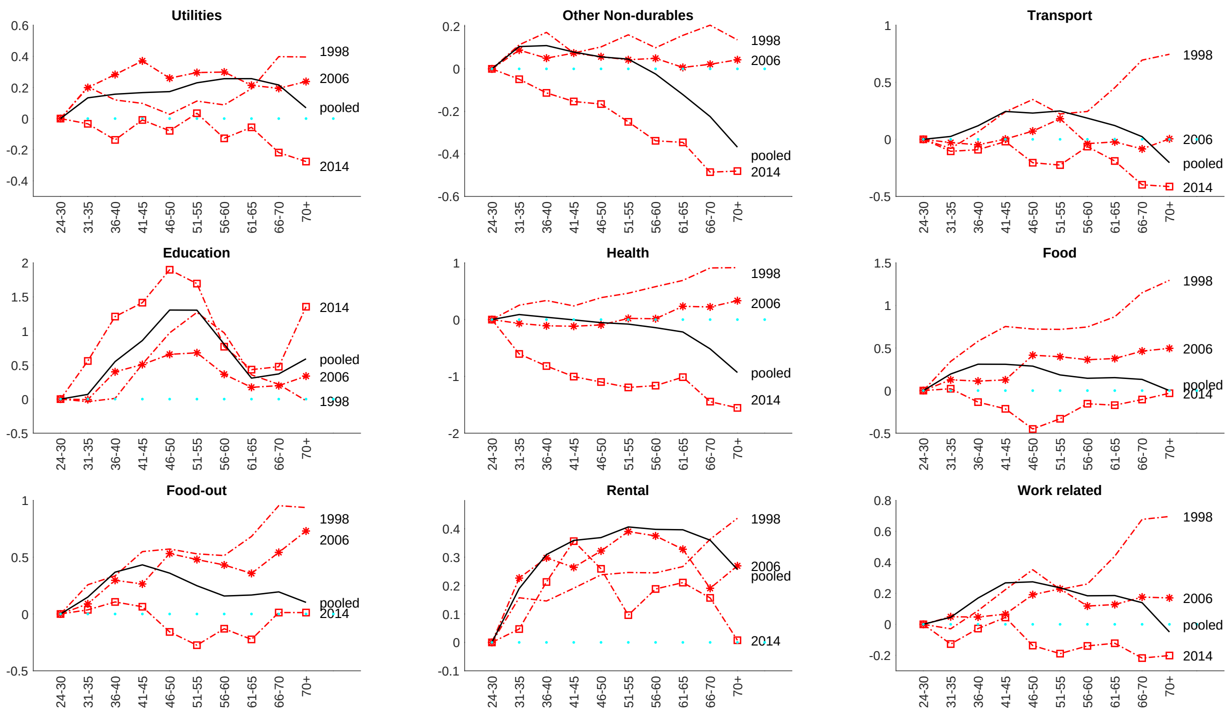

Figure A.13: Sub-categories age group coefficients: Each dashed line depicts $\beta_{g,t}$ for a selected year of the wave of the survey, $(1998, 2006$ and $2014)$, depicting the estimated lifecycle pattern of consumption for each year. The dark line depicts the age-group effects β_g when age-group effects pools information for the entire sample. The graph considers 10 age groups.

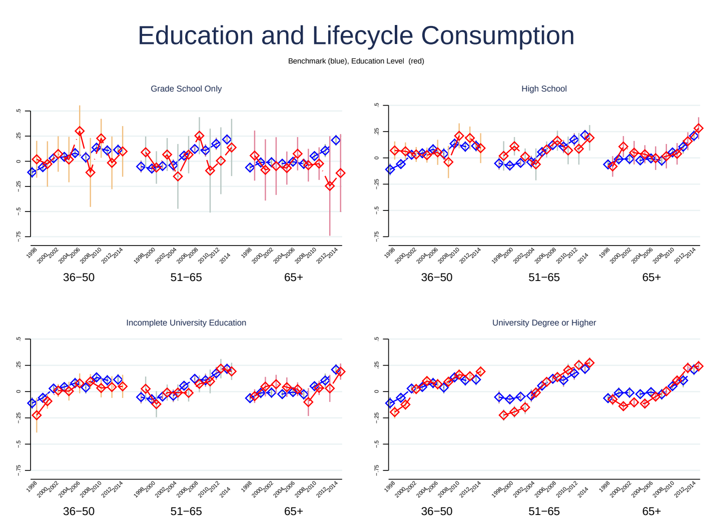

Figure A.14: Age Group Estimates for Groups with Different Education Levels (red) versus the Benchmark (blue)

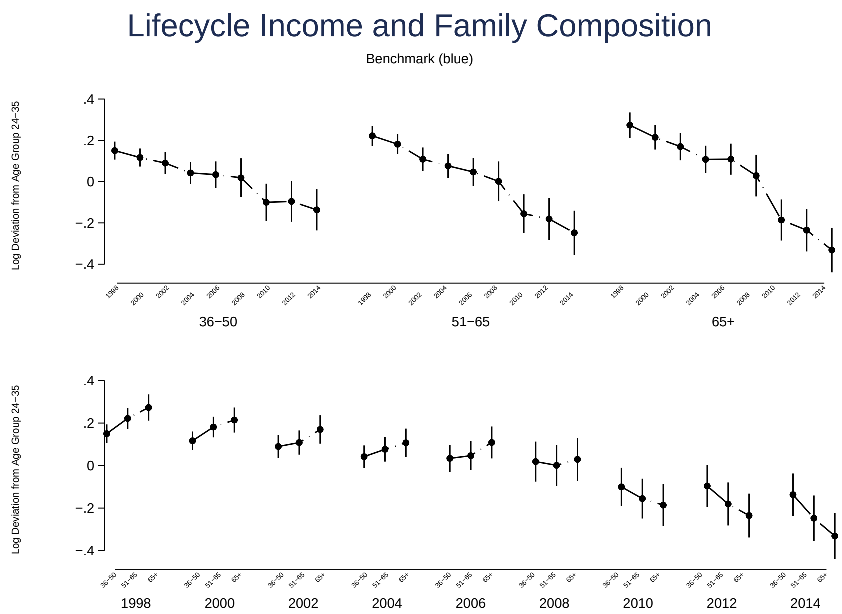

Figure A.15: Age Group Estimates including only Stable Households

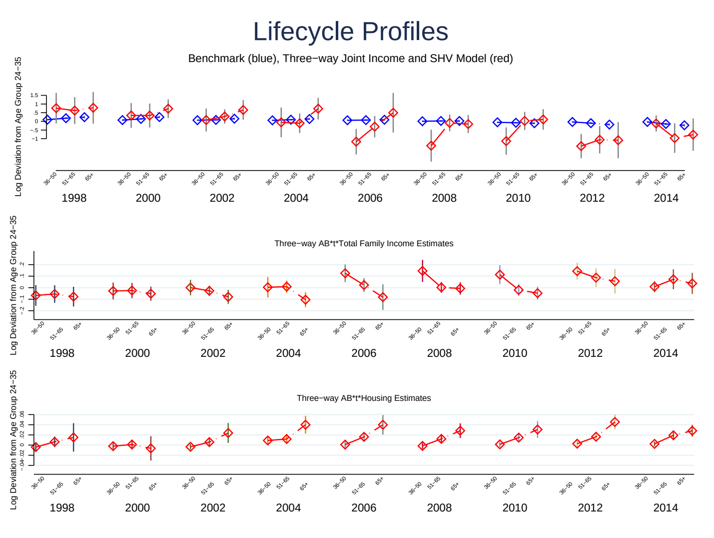

Figure A.16: Age group coefficients: Benchmark Model (Equation 10) and Joint Income and Housing Interaction Model (Equation 11), Whole Sample. Top Panel: $\beta_{g,t}$ from Equation 10 (blue) and Equation 11 (red); Middle Panel: $\theta_{g,Y,t}$ Bottom Panel: $\theta_{g,H,t}$

F Additional Results - Income Lifecycle Variation

We report the results of the income estimation using labour income instead of total family income (Figure A.17), by controlling for subjective house values (Figure A.18), for the imputed long sample (Figure A.20), and family composition with stable households (Figure A.21). Finally, Figure A.22 shows the effects of adding cohort effects. The flattening is less pronounced in the *Cohort View* than in the *Period View* model indicating that appropriately controlling for time effects instead of attributing trends to age and cohort and employing orthogonalised time dummies is relevant to uncover age time variation in income.

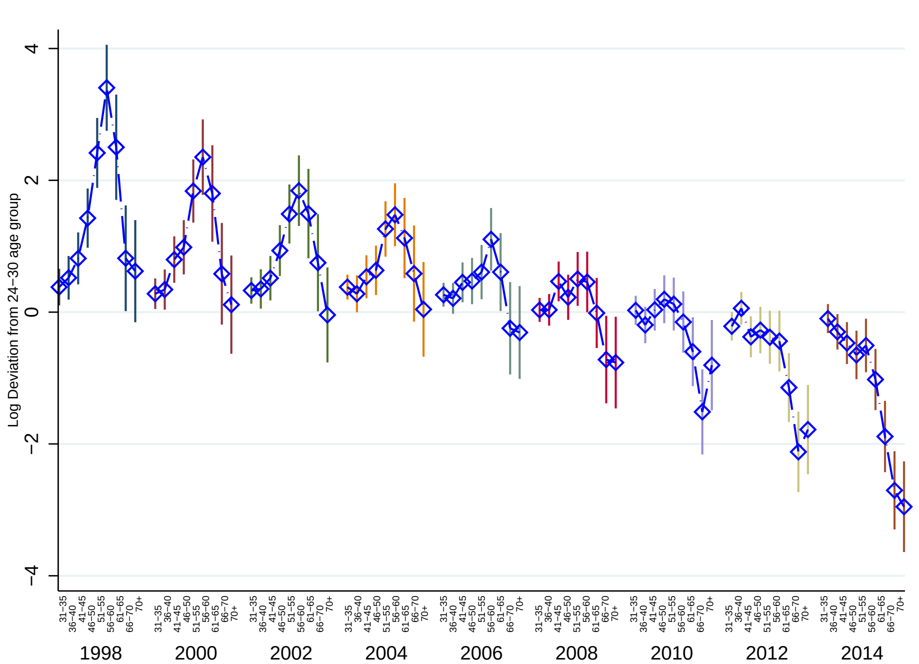

Figure A.17: Labour Income: Age group coefficients ($\beta^{Y}_{g,t}$) plotted by age group by year

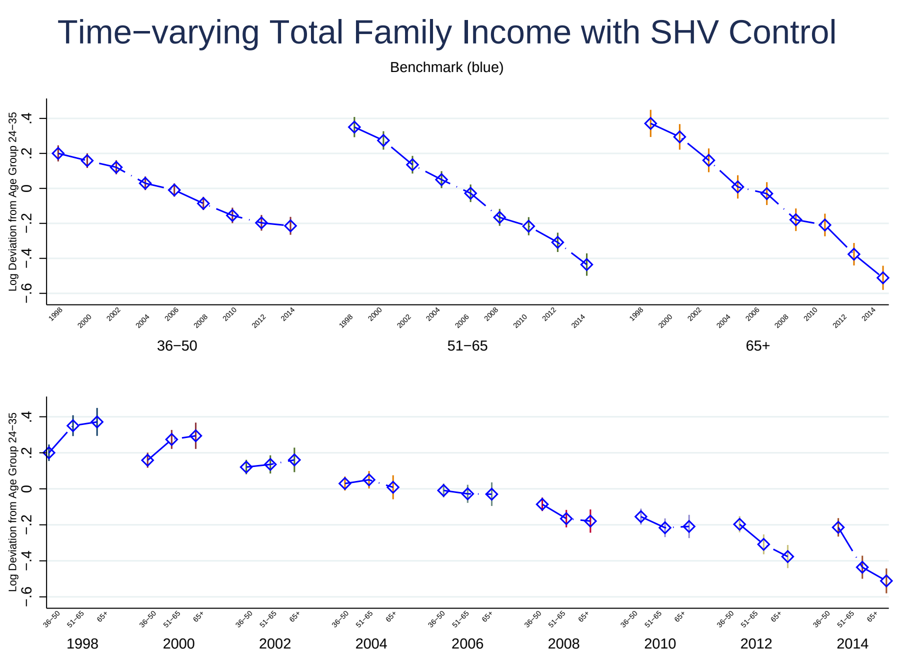

Figure A.18: Lifecycle Income: Age Group Estimates with Subjective House Value Control

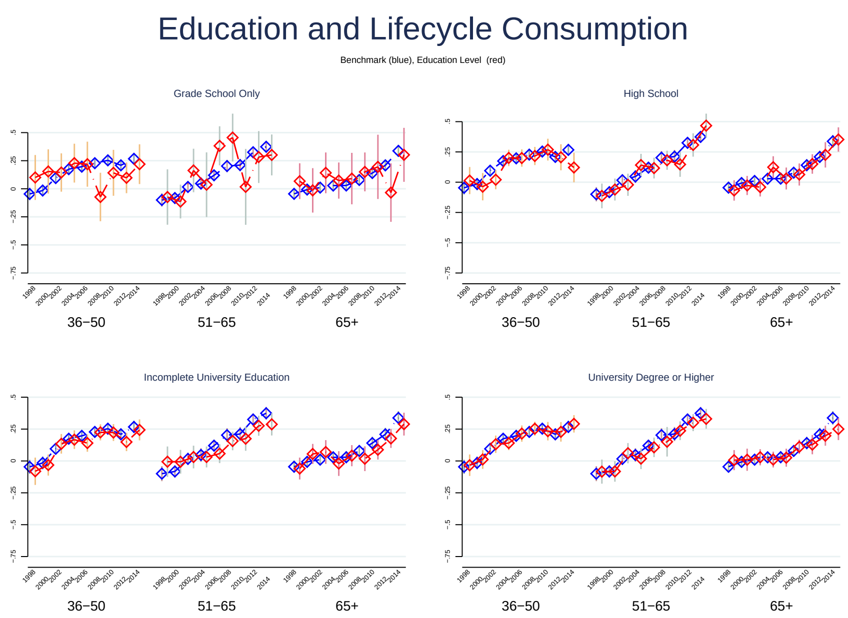

Figure A.19: Lifecycle Income: Age Group Estimates for Groups with Different Education Levels (red) versus the Benchmark (blue)

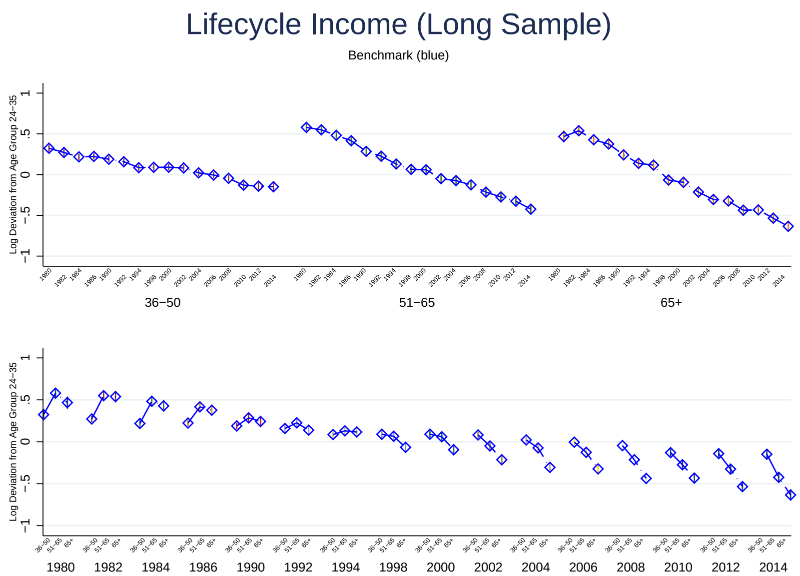

Figure A.20: Lifecycle Income: Age Group Estimates with Long (Imputed) Sample

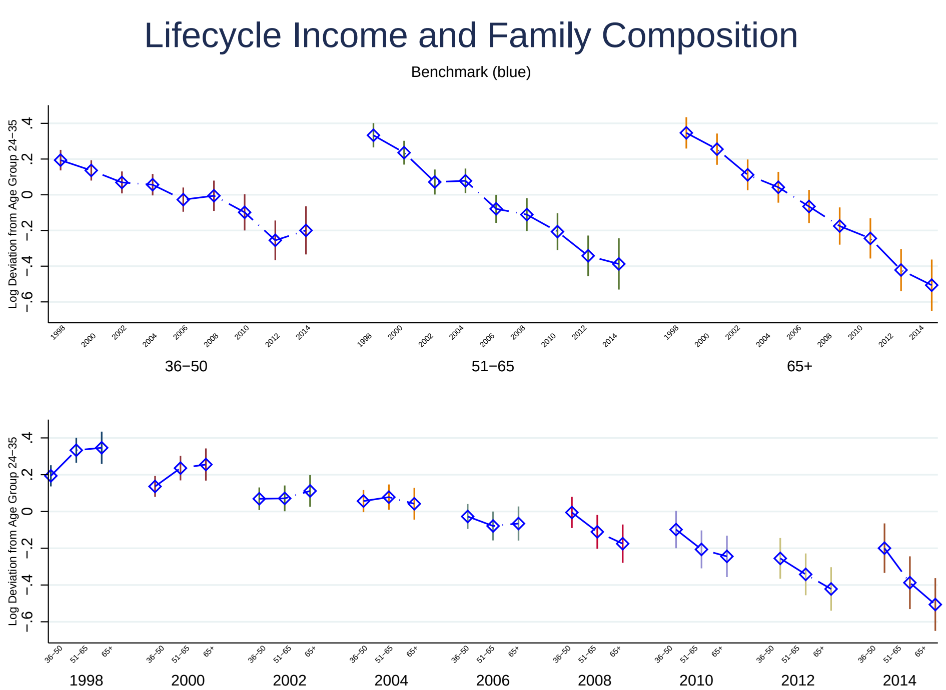

Figure A.21: Lifecycle Income: Age Group Estimates with Stable Households

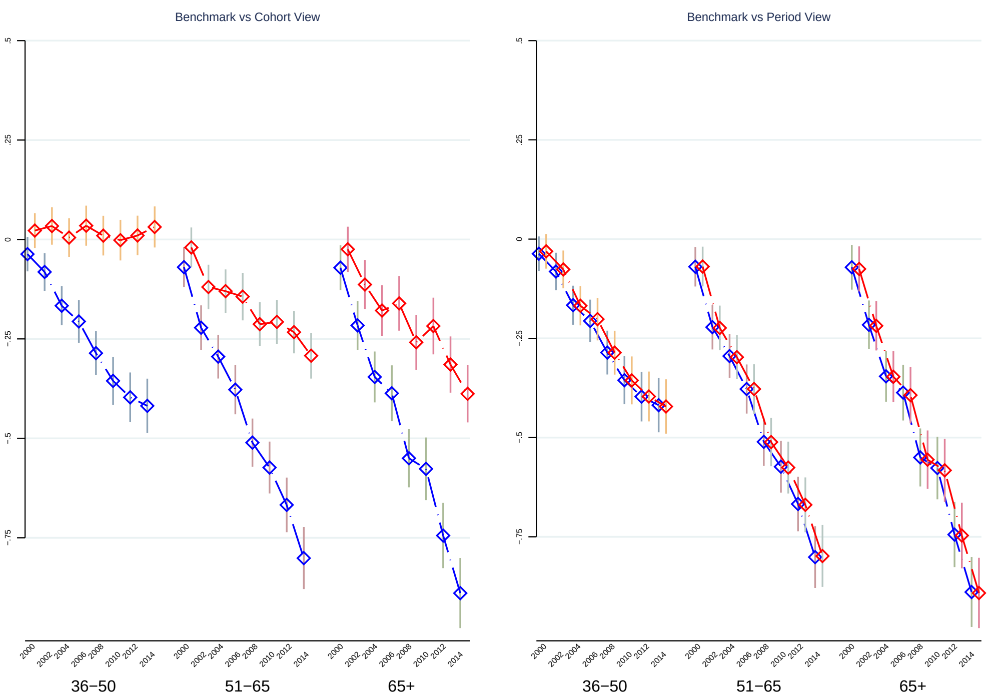

Figure A.22: Including Cohort Effects - Time-Varying Lifecycle Income
$\gamma^2_{g,t}$ - Benchmark (blue), Cohort Models (red)

G Estimation: Age-Time Effects with 4 Age Groups

Table A.5 show the full results of the estimation of the benchmark model Case 1 - *Pooled Lifecycle* and Case 2 - *Time-varying Lifecycle* with and without economic controls. Table A.6 show the full results of the estimation of the benchmark model Case 2 - *Time-varying Lifecycle* and the *Interaction Models*.

Table A.5: Benchmark Estimations: Nondurable Consumption Expenditures

	(1)		(2)		(3)	
	Pooled		Benchmark: Time-Varying Lifecycle		Benchmark: Time-varying Lifecycle with TFI and SHV	
AB2	0.08***	(0.00)				
AB3	0.07**	(0.00)				
AB4	-0.04	(0.25)				
t2000	0.05***	(0.00)	0.06**	(0.01)	0.04*	(0.05)
t2002	0.07***	(0.00)	0.10***	(0.00)	0.07**	(0.00)
t2004	0.07***	(0.00)	0.10***	(0.00)	0.05	(0.06)
t2006	0.05***	(0.00)	0.09***	(0.00)	0.03	(0.27)
t2008	-0.06***	(0.00)	0.03	(0.32)	-0.03	(0.24)
t2010	-0.02	(0.13)	0.13***	(0.00)	0.05	(0.07)
t2012	-0.07***	(0.00)	0.10**	(0.00)	0.02	(0.60)
t2014	-0.05**	(0.01)	0.14***	(0.00)	0.03	(0.36)
# Adults = 2	-0.08***	(0.00)	-0.09***	(0.00)	0.00	(0.82)
# Adults = 3	-0.20***	(0.00)	-0.22***	(0.00)	-0.05*	(0.04)
# Adults = 4	-0.32***	(0.00)	-0.34***	(0.00)	-0.09**	(0.00)
# Adults = 5	-0.45***	(0.00)	-0.48***	(0.00)	-0.18***	(0.00)
# Adults = 6	-0.46***	(0.00)	-0.47***	(0.00)	-0.10	(0.39)
# Adults = 7	-0.88***	(0.00)	-0.86***	(0.00)	-0.48**	(0.01)
# Adults = 8	0.57***	(0.00)	0.61***	(0.00)	0.73***	(0.00)
# Child = 2	-0.13***	(0.00)	-0.15***	(0.00)	-0.07***	(0.00)
# Child = 3	-0.29***	(0.00)	-0.31***	(0.00)	-0.17***	(0.00)
# Child = 4	-0.42***	(0.00)	-0.45***	(0.00)	-0.25***	(0.00)
# Child = 5	-0.55***	(0.00)	-0.59***	(0.00)	-0.34***	(0.00)
# Child = 6	-0.66***	(0.00)	-0.70***	(0.00)	-0.42***	(0.00)
# Child = 7	-0.51***	(0.00)	-0.55***	(0.00)	-0.22	(0.08)
# Child = 8	-0.52***	(0.00)	-0.58***	(0.00)	-0.23*	(0.04)
# Child = 9	-0.56***	(0.00)	-0.61***	(0.00)	-0.28**	(0.00)
# Child = 10	-0.74***	(0.00)	-0.79***	(0.00)	-0.53***	(0.00)
# Child = 11	-1.76***	(0.00)	-1.83***	(0.00)	-0.82***	(0.00)
# Child = 12	-0.80***	(0.00)	-0.95***	(0.00)	-0.61***	(0.00)
White	-0.09	(0.10)	-0.08	(0.10)	-0.08	(0.08)
Black	0.14	(0.18)	0.14	(0.20)	0.13	(0.21)
State 2	0.28**	(0.01)	0.27*	(0.01)	0.18	(0.08)
State 3	0.06	(0.71)	0.06	(0.71)	0.07	(0.69)
State 4	0.45***	(0.00)	0.43***	(0.00)	0.35**	(0.00)
State 5	0.35**	(0.00)	0.34**	(0.00)	0.25*	(0.02)
State 6	0.23	(0.11)	0.20	(0.16)	0.19	(0.17)
State 7	0.30*	(0.02)	0.29*	(0.02)	0.23	(0.08)
State 8	0.19	(0.13)	0.17	(0.17)	0.07	(0.59)
State 9	0.10	(0.34)	0.10	(0.37)	0.09	(0.39)
State 10	0.22*	(0.03)	0.21*	(0.04)	0.17	(0.10)
State 11	0.02	(0.89)	0.00	(1.00)	0.01	(0.95)
State 12	0.29*	(0.02)	0.28*	(0.03)	0.24	(0.06)
State 13	-0.08	(0.53)	-0.09	(0.46)	-0.15	(0.22)
State 14	0.07	(0.52)	0.07	(0.53)	0.07	(0.55)

State 15	0.09	(0.47)	0.08	(0.53)	0.09	(0.49)
State 16	-0.08	(0.56)	-0.10	(0.46)	-0.14	(0.31)
State 17	-0.21	(0.16)	-0.23	(0.13)	-0.20	(0.17)
State 18	0.39*	(0.03)	0.36*	(0.04)	0.19	(0.26)
State 19	0.27*	(0.02)	0.26*	(0.02)	0.14	(0.21)
State 20	0.17	(0.39)	0.16	(0.42)	0.26*	(0.05)
State 21	0.25	(0.12)	0.23	(0.15)	0.14	(0.30)
State 22	0.41**	(0.01)	0.39**	(0.01)	0.31*	(0.03)
State 23	0.30*	(0.02)	0.29*	(0.03)	0.25	(0.07)
State 24	0.10	(0.38)	0.09	(0.40)	0.09	(0.37)
State 25	0.36	(0.08)	0.35	(0.09)	0.23	(0.27)
State 26	0.06	(0.62)	0.05	(0.67)	-0.02	(0.89)
State 27	0.38***	(0.00)	0.36**	(0.00)	0.29*	(0.01)
State 28	0.44	(0.08)	0.44	(0.08)	0.35	(0.14)
State 29	0.25*	(0.04)	0.24*	(0.04)	0.16	(0.15)
State 30	0.29*	(0.02)	0.26*	(0.03)	0.14	(0.21)
State 31	0.29*	(0.01)	0.27*	(0.02)	0.21*	(0.05)
State 32	0.15	(0.14)	0.14	(0.19)	0.12	(0.24)
State 33	0.56***	(0.00)	0.57**	(0.00)	0.51***	(0.00)
State 34	0.18	(0.14)	0.17	(0.17)	0.12	(0.31)
State 35	-0.02	(0.91)	-0.03	(0.83)	-0.10	(0.46)
State 36	0.17	(0.18)	0.16	(0.22)	0.10	(0.41)
State 37	0.26*	(0.01)	0.25*	(0.02)	0.19	(0.06)
State 38	0.26	(0.26)	0.26	(0.26)	0.31	(0.17)
State 39	0.12	(0.29)	0.12	(0.31)	0.10	(0.39)
State 40	0.10	(0.46)	0.08	(0.57)	0.10	(0.43)
State 41	0.22	(0.09)	0.21	(0.10)	0.13	(0.32)
State 42	0.27*	(0.03)	0.26*	(0.04)	0.23	(0.05)
State 43	0.26*	(0.04)	0.25	(0.05)	0.16	(0.18)
State 44	0.01	(0.98)	-0.03	(0.87)	-0.10	(0.55)
State 45	0.21	(0.06)	0.20	(0.08)	0.12	(0.30)
State 46	0.23	(0.05)	0.23	(0.06)	0.19	(0.10)
State 47	0.44**	(0.01)	0.43**	(0.01)	0.36*	(0.02)
State 48	0.19	(0.13)	0.16	(0.19)	0.08	(0.53)
State 49	0.27*	(0.03)	0.26*	(0.04)	0.13	(0.30)
State 50	0.42*	(0.01)	0.39*	(0.02)	0.29	(0.07)
State 51	0.66***	(0.00)	0.67***	(0.00)	0.55***	(0.00)
Nohome	-0.24***	(0.00)	-0.23***	(0.00)	1.85***	(0.00)
SelfEmp	0.03	(0.11)	0.02	(0.25)	0.00	(0.79)
Disability	-0.06***	(0.00)	-0.06***	(0.00)	-0.03*	(0.04)
Marital_Status	0.00	(0.83)	-0.01	(0.58)	0.00	(1.00)
AB2xt1998			0.10**	(0.00)	0.06*	(0.04)
AB2xt2000			0.07*	(0.01)	0.03	(0.27)
AB2xt2002			0.07**	(0.01)	0.04	(0.09)
AB2xt2004			0.06*	(0.02)	0.04	(0.14)
AB2xt2006			0.07**	(0.01)	0.04	(0.11)
AB2xt2008			0.01	(0.75)	-0.01	(0.75)
AB2xt2010			-0.05	(0.06)	-0.04	(0.10)
AB2xt2012			-0.03	(0.28)	-0.03	(0.30)
AB2xt2014			-0.03	(0.35)	-0.01	(0.67)
AB3xt1998			0.18***	(0.00)	0.11***	(0.00)
AB3xt2000			0.14***	(0.00)	0.07*	(0.02)

AB3xt2002		0.08**	(0.00)	0.04	(0.14)
AB3xt2004		0.10**	(0.00)	0.07*	(0.03)
AB3xt2006		0.08*	(0.01)	0.05	(0.10)
AB3xt2008		0.03	(0.43)	0.03	(0.36)
AB3xt2010		-0.08*	(0.02)	-0.05	(0.09)
AB3xt2012		-0.10**	(0.01)	-0.07*	(0.03)
AB3xt2014		-0.15***	(0.00)	-0.09*	(0.02)
AB4xt1998		0.24***	(0.00)	0.15***	(0.00)
AB4xt2000		0.25***	(0.00)	0.15***	(0.00)
AB4xt2002		0.15***	(0.00)	0.09*	(0.01)
AB4xt2004		0.14***	(0.00)	0.10**	(0.01)
AB4xt2006		0.09*	(0.02)	0.06	(0.10)
AB4xt2008		0.02	(0.53)	0.02	(0.48)
AB4xt2010		-0.11**	(0.01)	-0.08*	(0.03)
AB4xt2012		-0.18***	(0.00)	-0.13***	(0.00)
AB4xt2014		-0.22***	(0.00)	-0.15***	(0.00)
Income				0.17***	(0.00)
SHV				0.19***	(0.00)
Observations	44149	44149		43512	
Adjusted R^2	0.058	0.062		0.115	
AIC	56654.56	56465.81		49054.34	
BIC	57367.57	57387.51		49991.86	

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.6: Nondurable Consumption Expenditures

	(1) Benchmark		(2) Interaction: SHV		(3) Interaction: TFI		(4) Interaction: Joint SHV and TFI	
AB2xt1998	0.06*	(0.04)	0.19***	(0.00)	-0.24	(0.63)	-0.48	(0.35)
AB2xt2000	0.03	(0.27)	0.11*	(0.03)	-0.76	(0.07)	-0.95*	(0.02)
AB2xt2002	0.04	(0.09)	0.08	(0.07)	-1.08**	(0.00)	-1.21**	(0.00)
AB2xt2004	0.04	(0.14)	-0.07	(0.25)	-1.42**	(0.00)	-1.38**	(0.01)
AB2xt2006	0.04	(0.11)	-0.07	(0.11)	-2.33***	(0.00)	-2.41***	(0.00)
AB2xt2008	-0.01	(0.75)	-0.12*	(0.02)	-2.33***	(0.00)	-2.61***	(0.00)
AB2xt2010	-0.04	(0.10)	-0.17***	(0.00)	-2.17***	(0.00)	-2.28***	(0.00)
AB2xt2012	-0.03	(0.30)	-0.17***	(0.00)	-2.49***	(0.00)	-2.56***	(0.00)
AB2xt2014	-0.01	(0.67)	-0.11*	(0.02)	-1.25***	(0.00)	-1.20***	(0.00)
AB3xt1998	0.11***	(0.00)	0.15	(0.06)	-0.50	(0.26)	-0.74	(0.10)
AB3xt2000	0.07*	(0.02)	0.13	(0.08)	-0.76	(0.06)	-0.99*	(0.01)
AB3xt2002	0.04	(0.14)	0.03	(0.64)	-0.90***	(0.00)	-1.12***	(0.00)
AB3xt2004	0.07*	(0.03)	-0.08	(0.31)	-1.45***	(0.00)	-1.53***	(0.00)
AB3xt2006	0.05	(0.10)	-0.17*	(0.01)	-1.69***	(0.00)	-1.78***	(0.00)
AB3xt2008	0.03	(0.36)	-0.15*	(0.03)	-1.41***	(0.00)	-1.53***	(0.00)
AB3xt2010	-0.05	(0.09)	-0.25***	(0.00)	-1.44***	(0.00)	-1.47***	(0.00)
AB3xt2012	-0.07*	(0.03)	-0.36***	(0.00)	-2.56***	(0.00)	-2.56***	(0.00)
AB3xt2014	-0.09*	(0.02)	-0.42***	(0.00)	-2.51***	(0.00)	-2.44***	(0.00)
AB4xt1998	0.15***	(0.00)	0.03	(0.86)	-0.64	(0.27)	-0.70	(0.19)
AB4xt2000	0.15***	(0.00)	0.28	(0.07)	-0.40	(0.20)	-0.71**	(0.01)
AB4xt2002	0.09*	(0.01)	-0.10	(0.42)	-0.65*	(0.04)	-0.85**	(0.00)
AB4xt2004	0.10**	(0.01)	-0.24*	(0.05)	-0.85*	(0.02)	-0.81*	(0.02)
AB4xt2006	0.06	(0.10)	-0.31*	(0.05)	-1.11	(0.10)	-1.01	(0.14)
AB4xt2008	0.02	(0.48)	-0.26**	(0.01)	-1.60***	(0.00)	-1.66***	(0.00)
AB4xt2010	-0.08*	(0.03)	-0.42***	(0.00)	-1.24***	(0.00)	-1.37***	(0.00)
AB4xt2012	-0.13***	(0.00)	-0.79***	(0.00)	-2.62***	(0.00)	-2.59***	(0.00)
AB4xt2014	-0.15***	(0.00)	-0.58***	(0.00)	-2.32***	(0.00)	-2.29***	(0.00)
t2000	0.04*	(0.05)	0.07**	(0.00)	0.06**	(0.01)	0.07**	(0.00)
t2002	0.07**	(0.00)	0.12***	(0.00)	0.09***	(0.00)	0.11***	(0.00)
t2004	0.05	(0.06)	0.12***	(0.00)	0.10***	(0.00)	0.12***	(0.00)
t2006	0.03	(0.27)	0.12***	(0.00)	0.09***	(0.00)	0.12***	(0.00)
t2008	-0.03	(0.24)	0.06	(0.05)	0.04	(0.10)	0.07*	(0.01)
t2010	0.05	(0.07)	0.16***	(0.00)	0.13***	(0.00)	0.15***	(0.00)
t2012	0.02	(0.60)	0.13***	(0.00)	0.11***	(0.00)	0.13***	(0.00)
t2014	0.03	(0.36)	0.19***	(0.00)	0.14***	(0.00)	0.18***	(0.00)
# Adults = 2	0.00	(0.82)	-0.08***	(0.00)	-0.09***	(0.00)	-0.07***	(0.00)
# Adults = 3	-0.05*	(0.04)	-0.21***	(0.00)	-0.19***	(0.00)	-0.18***	(0.00)
# Adults = 4	-0.09**	(0.00)	-0.32***	(0.00)	-0.29***	(0.00)	-0.27***	(0.00)
# Adults = 5	-0.18***	(0.00)	-0.44***	(0.00)	-0.40***	(0.00)	-0.38***	(0.00)
# Adults = 6	-0.10	(0.39)	-0.36**	(0.00)	-0.42***	(0.00)	-0.31**	(0.01)
# Adults = 7	-0.48**	(0.01)	-0.82***	(0.00)	-0.68***	(0.00)	-0.68***	(0.00)
# Adults = 8	0.73***	(0.00)	0.43***	(0.00)	0.71***	(0.00)	0.46***	(0.00)
# Child = 2	-0.07***	(0.00)	-0.14***	(0.00)	-0.13***	(0.00)	-0.12***	(0.00)
# Child = 3	-0.17***	(0.00)	-0.30***	(0.00)	-0.28***	(0.00)	-0.27***	(0.00)
# Child = 4	-0.25***	(0.00)	-0.44***	(0.00)	-0.41***	(0.00)	-0.40***	(0.00)
# Child = 5	-0.34***	(0.00)	-0.58***	(0.00)	-0.53***	(0.00)	-0.52***	(0.00)
# Child = 6	-0.42***	(0.00)	-0.69***	(0.00)	-0.63***	(0.00)	-0.62***	(0.00)
# Child = 7	-0.22	(0.08)	-0.53***	(0.00)	-0.43**	(0.00)	-0.41**	(0.00)
# Child = 8	-0.23*	(0.04)	-0.59***	(0.00)	-0.47***	(0.00)	-0.46***	(0.00)

# Child = 9	-0.28**	(0.00)	-0.58***	(0.00)	-0.42***	(0.00)	-0.42***	(0.00)
# Child = 10	-0.53***	(0.00)	-0.77***	(0.00)	-0.56***	(0.00)	-0.55***	(0.00)
# Child = 11	-0.82***	(0.00)	-1.74***	(0.00)	-1.57***	(0.00)	-1.52***	(0.00)
# Child = 12	-0.61***	(0.00)	-0.99***	(0.00)	-0.93***	(0.00)	-0.94***	(0.00)
White	-0.08	(0.08)	-0.07	(0.18)	-0.07	(0.11)	-0.06	(0.20)
Black	0.13	(0.21)	0.17	(0.11)	0.12	(0.22)	0.15	(0.14)
State 2	0.18	(0.08)	0.25*	(0.02)	0.25*	(0.02)	0.23*	(0.03)
State 3	0.07	(0.69)	0.05	(0.79)	0.09	(0.59)	0.08	(0.66)
State 4	0.35**	(0.00)	0.41***	(0.00)	0.41***	(0.00)	0.39***	(0.00)
State 5	0.25*	(0.02)	0.32**	(0.00)	0.31**	(0.00)	0.30**	(0.01)
State 6	0.19	(0.17)	0.21	(0.14)	0.24	(0.07)	0.26	(0.06)
State 7	0.23	(0.08)	0.28*	(0.03)	0.29*	(0.02)	0.28*	(0.02)
State 8	0.07	(0.59)	0.15	(0.22)	0.13	(0.29)	0.12	(0.35)
State 9	0.09	(0.39)	0.09	(0.38)	0.08	(0.43)	0.07	(0.50)
State 10	0.17	(0.10)	0.21*	(0.04)	0.18	(0.08)	0.17	(0.09)
State 11	0.01	(0.95)	-0.03	(0.86)	-0.02	(0.92)	-0.04	(0.84)
State 12	0.24	(0.06)	0.27*	(0.03)	0.30*	(0.02)	0.29*	(0.02)
State 13	-0.15	(0.22)	-0.10	(0.43)	-0.11	(0.39)	-0.12	(0.32)
State 14	0.07	(0.55)	0.07	(0.53)	0.05	(0.65)	0.04	(0.69)
State 15	0.09	(0.49)	0.09	(0.48)	0.06	(0.67)	0.05	(0.68)
State 16	-0.14	(0.31)	-0.13	(0.37)	-0.17	(0.23)	-0.19	(0.19)
State 17	-0.20	(0.17)	-0.22	(0.14)	-0.22	(0.13)	-0.23	(0.13)
State 18	0.19	(0.26)	0.31	(0.10)	0.31	(0.08)	0.26	(0.17)
State 19	0.14	(0.21)	0.25*	(0.03)	0.20	(0.06)	0.19	(0.08)
State 20	0.26*	(0.05)	0.17	(0.40)	0.31*	(0.01)	0.32*	(0.01)
State 21	0.14	(0.30)	0.21	(0.21)	0.18	(0.18)	0.16	(0.25)
State 22	0.31*	(0.03)	0.37*	(0.01)	0.36*	(0.01)	0.34*	(0.02)
State 23	0.25	(0.07)	0.30*	(0.03)	0.30*	(0.02)	0.30*	(0.03)
State 24	0.09	(0.37)	0.09	(0.44)	0.09	(0.40)	0.08	(0.44)
State 25	0.23	(0.27)	0.37	(0.10)	0.35	(0.08)	0.33	(0.13)
State 26	-0.02	(0.89)	0.05	(0.69)	0.06	(0.58)	0.05	(0.66)
State 27	0.29*	(0.01)	0.36**	(0.00)	0.34**	(0.00)	0.34**	(0.00)
State 28	0.35	(0.14)	0.41	(0.09)	0.43	(0.06)	0.43	(0.06)
State 29	0.16	(0.15)	0.24*	(0.04)	0.21	(0.06)	0.21	(0.05)
State 30	0.14	(0.21)	0.23	(0.06)	0.21	(0.08)	0.19	(0.11)
State 31	0.21*	(0.05)	0.29*	(0.01)	0.21*	(0.04)	0.21*	(0.04)
State 32	0.12	(0.24)	0.13	(0.20)	0.10	(0.32)	0.10	(0.34)
State 33	0.51***	(0.00)	0.52**	(0.00)	0.57**	(0.00)	0.50**	(0.00)
State 34	0.12	(0.31)	0.16	(0.19)	0.16	(0.18)	0.15	(0.22)
State 35	-0.10	(0.46)	-0.06	(0.64)	-0.04	(0.77)	-0.06	(0.64)
State 36	0.10	(0.41)	0.13	(0.30)	0.15	(0.23)	0.13	(0.30)
State 37	0.19	(0.06)	0.24*	(0.02)	0.24*	(0.02)	0.23*	(0.03)
State 38	0.31	(0.17)	0.20	(0.38)	0.35	(0.16)	0.31	(0.21)
State 39	0.10	(0.39)	0.11	(0.31)	0.08	(0.49)	0.09	(0.44)
State 40	0.10	(0.43)	0.13	(0.36)	0.08	(0.55)	0.12	(0.42)
State 41	0.13	(0.32)	0.19	(0.15)	0.16	(0.20)	0.15	(0.24)
State 42	0.23	(0.05)	0.26*	(0.03)	0.23*	(0.04)	0.23*	(0.04)
State 43	0.16	(0.18)	0.22	(0.08)	0.21	(0.09)	0.18	(0.13)
State 44	-0.10	(0.55)	-0.03	(0.87)	-0.05	(0.76)	-0.03	(0.87)
State 45	0.12	(0.30)	0.21	(0.06)	0.18	(0.10)	0.19	(0.09)
State 46	0.19	(0.10)	0.21	(0.08)	0.22	(0.06)	0.21	(0.07)
State 47	0.36*	(0.02)	0.37*	(0.01)	0.51**	(0.00)	0.43**	(0.00)

State 48	0.08	(0.53)	0.14	(0.24)	0.12	(0.31)	0.10	(0.42)
State 49	0.13	(0.30)	0.22	(0.09)	0.21	(0.08)	0.19	(0.15)
State 50	0.29	(0.07)	0.41^{**}	(0.01)	0.34^{*}	(0.03)	0.37^{*}	(0.02)
State 51	0.55^{***}	(0.00)	0.64^{***}	(0.00)	0.64^{***}	(0.00)	0.61^{***}	(0.00)
Nohome	1.85^{***}	(0.00)	-0.11^{***}	(0.00)	-0.21^{***}	(0.00)	-0.12^{***}	(0.00)
SelfEmp	0.00	(0.79)	0.03	(0.11)	0.00	(0.79)	0.01	(0.57)
Disability	-0.03^{*}	(0.04)	-0.05^{***}	(0.00)	-0.03^{*}	(0.04)	-0.03	(0.06)
Marital_Status	0.00	(1.00)	-0.00	(0.86)	-0.00	(0.64)	-0.00	(0.78)
Income	0.17^{***}	(0.00)						
SHV	0.19^{***}	(0.00)						
SHVxAB2xt1998			-0.01	(0.12)			-0.00	(0.57)
SHVxAB2xt2000			-0.00	(0.67)			0.00	(0.98)
SHVxAB2xt2002			0.00	(0.74)			0.00	(0.93)
SHVxAB2xt2004			0.02^{***}	(0.00)			0.01^{**}	(0.00)
SHVxAB2xt2006			0.02^{***}	(0.00)			0.01	(0.14)
SHVxAB2xt2008			0.02^{***}	(0.00)			0.00	(0.89)
SHVxAB2xt2010			0.02^{***}	(0.00)			0.00	(0.47)
SHVxAB2xt2012			0.02^{***}	(0.00)			0.00	(0.41)
SHVxAB2xt2014			0.01^{*}	(0.02)			0.00	(0.39)
SHVxAB3xt1998			0.01	(0.21)			0.01	(0.09)
SHVxAB3xt2000			0.01	(0.34)			0.01	(0.29)
SHVxAB3xt2002			0.01^{*}	(0.03)			0.01^{*}	(0.03)
SHVxAB3xt2004			0.02^{***}	(0.00)			0.02^{**}	(0.00)
SHVxAB3xt2006			0.03^{***}	(0.00)			0.02^{***}	(0.00)
SHVxAB3xt2008			0.02^{***}	(0.00)			0.02^{***}	(0.00)
SHVxAB3xt2010			0.02^{***}	(0.00)			0.02^{***}	(0.00)
SHVxAB3xt2012			0.03^{***}	(0.00)			0.02^{***}	(0.00)
SHVxAB3xt2014			0.03^{***}	(0.00)			0.02^{***}	(0.00)
SHVxAB4xt1998			0.02	(0.19)			0.02	(0.29)
SHVxAB4xt2000			-0.00	(0.81)			-0.00	(0.92)
SHVxAB4xt2002			0.03^{*}	(0.01)			0.03^{*}	(0.02)
SHVxAB4xt2004			0.04^{***}	(0.00)			0.05^{***}	(0.00)
SHVxAB4xt2006			0.05^{**}	(0.00)			0.05^{***}	(0.00)
SHVxAB4xt2008			0.03^{***}	(0.00)			0.03^{***}	(0.00)
SHVxAB4xt2010			0.04^{***}	(0.00)			0.04^{***}	(0.00)
SHVxAB4xt2012			0.07^{***}	(0.00)			0.05^{***}	(0.00)
SHVxAB4xt2014			0.04^{***}	(0.00)			0.03^{***}	(0.00)
IncomexAB2xt1998					0.04	(0.46)	0.07	(0.21)
IncomexAB2xt2000					0.09^{*}	(0.03)	0.11^{*}	(0.01)
IncomexAB2xt2002					0.12^{**}	(0.00)	0.14^{***}	(0.00)
IncomexAB2xt2004					0.16^{**}	(0.00)	0.14^{**}	(0.01)
IncomexAB2xt2006					0.25^{***}	(0.00)	0.25^{***}	(0.00)
IncomexAB2xt2008					0.24^{***}	(0.00)	0.27^{***}	(0.00)
IncomexAB2xt2010					0.22^{***}	(0.00)	0.23^{***}	(0.00)
IncomexAB2xt2012					0.26^{***}	(0.00)	0.26^{***}	(0.00)
IncomexAB2xt2014					0.13^{***}	(0.00)	0.12^{***}	(0.00)
IncomexAB3xt1998					0.07	(0.11)	0.09^{*}	(0.04)
IncomexAB3xt2000					0.09^{*}	(0.02)	0.11^{**}	(0.00)
IncomexAB3xt2002					0.10^{***}	(0.00)	0.12^{***}	(0.00)
IncomexAB3xt2004					0.16^{***}	(0.00)	0.15^{***}	(0.00)
IncomexAB3xt2006					0.18^{***}	(0.00)	0.17^{***}	(0.00)
IncomexAB3xt2008					0.15^{***}	(0.00)	0.15^{***}	(0.00)

IncomexAB3xt2010			0.14***	(0.00)	0.13***	(0.00)
IncomexAB3xt2012			0.25***	(0.00)	0.24***	(0.00)
IncomexAB3xt2014			0.24***	(0.00)	0.22***	(0.00)
IncomexAB4xt1998			0.09	(0.13)	0.08	(0.10)
IncomexAB4xt2000			0.06*	(0.04)	0.10**	(0.00)
IncomexAB4xt2002			0.08**	(0.01)	0.08*	(0.02)
IncomexAB4xt2004			0.10**	(0.00)	0.05	(0.12)
IncomexAB4xt2006			0.13	(0.07)	0.07	(0.29)
IncomexAB4xt2008			0.17***	(0.00)	0.15***	(0.00)
IncomexAB4xt2010			0.12***	(0.00)	0.10***	(0.00)
IncomexAB4xt2012			0.26***	(0.00)	0.21***	(0.00)
IncomexAB4xt2014			0.22***	(0.00)	0.19***	(0.00)
Observations	43512	43641	44010		43512	
Adjusted R^2	0.115	0.069	0.099		0.101	
AIC	49054.34	55150.80	50870.32		49824.91	
BIC	49991.86	56305.74	52026.38		51213.84	

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

H Asset Profiles - Theoretical model

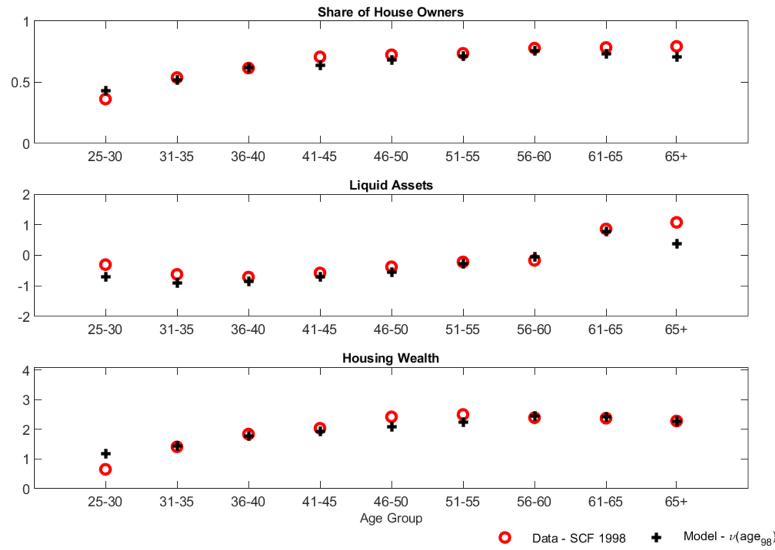

Figure A.23: Calibration: Lifecycle profiles of homeownership, non-housing wealth and housing wealth. Model versus SCF 1998 Data

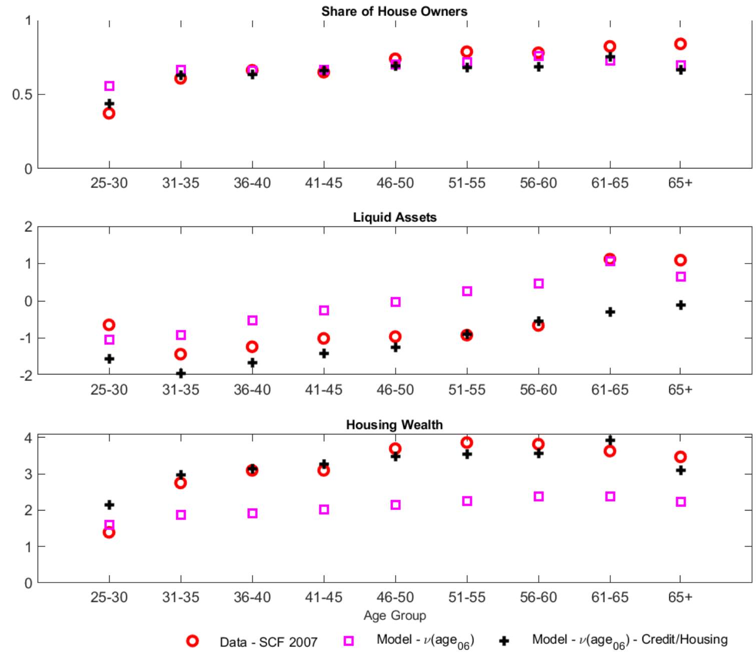

Figure A.24: Asset Life Cycle Profiles: The Role of Changes in Income Profiles and Credit/House Prices During the Boom. Data: Survey of Consumer Finance - 2007)

Note: Model $\nu(age_{06})$, incorporates the change in income profile only (pink squares) - Model $\nu(age_{06})$ - Credit/Housing, incorporates change in income profile, the relaxation in credit constraints and the increase in the trend in house prices (dark plus sign).

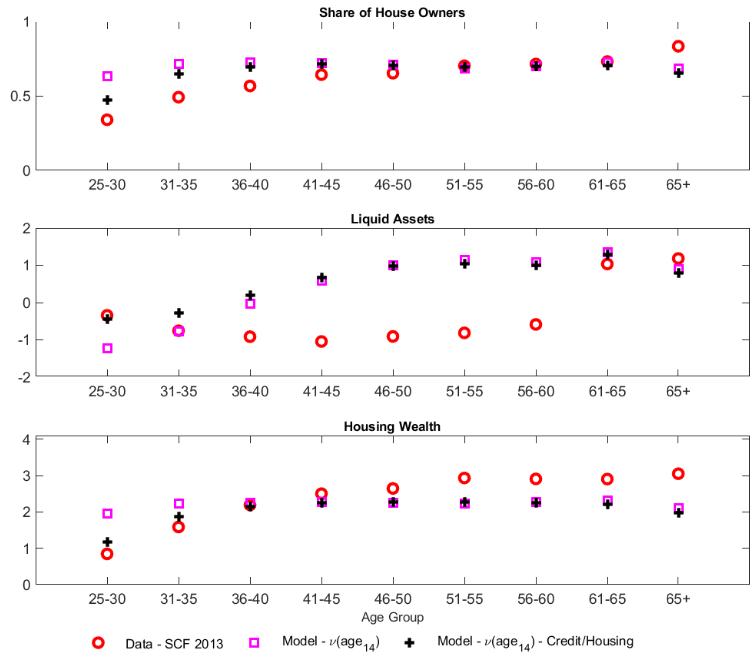

Figure A.25: Asset Life Cycle Profiles: The Role of Changes in Income Profiles and Credit/House Prices During the Recession. Data: Survey of Consumer Finance - 2013)

Note: Model $\nu(age_{14})$, incorporates the change in income profile only (pink squares) - Model $\nu(age_{14})$ - Credit/Housing, incorporates change in income profile, the tightening in credit constraints and the decrease in the trend in house prices (dark plus sign).