Demographics and Monetary Policy Shocks*

Kimberly A. Berg[†]
Miami University

Chadwick C. Curtis[‡]
University of Richmond

Steven Lugauer[§]
University of Kentucky

Nelson C. Mark[¶]
University of Notre Dame and NBER

This Draft: September 7, 2018

Abstract

We study consumption responses of households at different stages of the life cycle to monetary policy shocks. Older households have a higher consumption response than younger households. Amongst older households, the consumption response is increasing with income, suggesting that wealth effects play a role in driving the response patterns. A partial-equilibrium life-cycle model of consumption-saving and labor-supply decisions can qualitatively explain these empirical patterns. Understanding the heterogeneity in consumption responses across age groups is important for understanding the transmission of monetary policy, especially as the U.S. population grows older.

Keywords: Monetary Policy Shocks, Consumption, Demographic Change, Life-Cycle
JEL Classification: E4, E52, E21, J11, D15

---

*We thank Ben Pugsley for comments on an earlier draft and seminar participants at the Robins School of Business, Miami University, and the Liberal Arts Macro Conference.

[†]E-mail address: bergka@miamioh.edu.
[‡]E-mail address: ccurtis2@richmond.edu.
[§]E-mail address: steven.lugauer@uky.edu.
[¶]E-mail address: nmark@nd.edu.

1 Introduction

This paper studies the responses of household consumption expenditures across age groups to monetary policy shocks. Empirically, we find that monetary policy shocks have a larger impact on consumption expenditures for older households. That is, the interest rate semi-elasticity of consumption expenditures is higher for the old, relative to the young. Within the older age group, we also find the consumption response to monetary policy shocks to be increasing in income. Since income is correlated with wealth, and older households have higher net-wealth than younger households, this evidence suggests that a wealth-effect plays a significant role in driving the differential consumption response patterns. To better understand the mechanism that underlies these empirical patterns, we study a partial-equilibrium life-cycle model of consumption, saving, and labor-supply decisions. The model endogenously produces age-related consumption response heterogeneity to interest rate shocks in a manner that is largely consistent with the data.

Figure 1: Ratio of U.S. Population Aged 65+ to Population Aged 25-64: 1950-2050

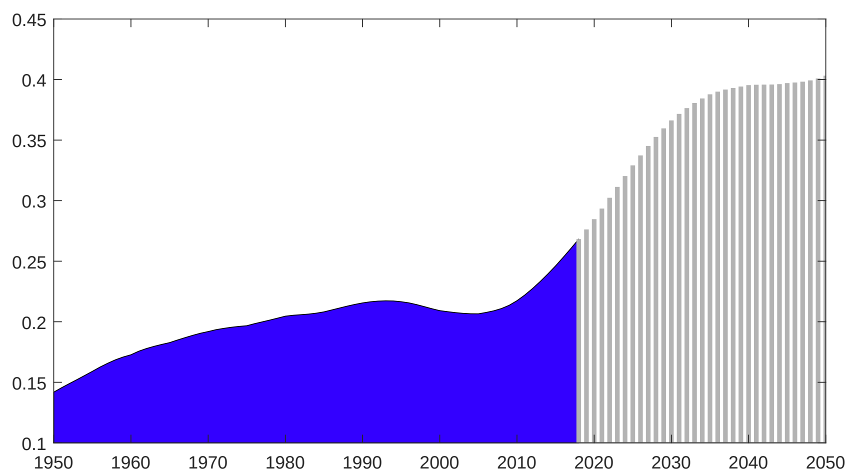

Notes: Data is from the UN World Population Prospects 2017 Revision. The gray area are projected figures.

Two motivations drive our inquiry. First, since consumption is the largest component of GDP, a better understanding of age related heterogeneity in consumption responses to monetary policy shocks can improve our knowledge of the aggregate transmission channel and about those population segments most impacted by monetary policy. The second motivation is the accelerating demographic transition towards an older population, currently underway in the U.S. and other developed economies. As seen in Figure 1, the ratio of the U.S. population over 65 to those between 25 and 64 is rapidly (in demographic time) increasing. The ratio, which had been fairly steady around 0.2 from 1980 to 2010, is projected to double as post-war baby boomers age into retirement. Age-related heterogeneity in consumption responses could potentially change the effectiveness of monetary policy as the population ages.

Our empirical analysis is based on impulse responses from a structural vector autoregression (VAR) and on local projections (Jordà (2005)). Employing consumption data from the Consumer Expenditure Survey (CEX), we classify households into young (household head aged 25-34), middle (35-64), and old (65+), and study how their consumption responds to three alternative monetary policy shocks that have been constructed by other researchers. The shock series we employ are constructed using the high-frequency identification methodology of Barakchian and Crowe (2013) and Gürkaynak et al. [2005], and the narrative/Greenbook methodology of Romer and Romer (2004). Our general finding across two empirical methods and three identified shock series is that old households have the highest proportionate consumption response to monetary policy shocks. Middle-aged household consumption appears to be more responsive than young households, but the evidence here is less definitive.

The interest rate shock exerts a pure substitution effect, but beyond this, we conjecture four life-cycle related effects at work in driving the observed age-related consumption patterns. First, older households tend to be wealthier than younger households, so a given decline in the interest rate generates a larger capital gain for the old. Second, older households may be more sensitive to interest rate changes due to the composition of their portfolios. Our analysis of data from the Survey of Consumer Finances (SCF) shows that the composition of wealth for older households is tilted towards long-term assets (home equity, bond retirement funds, and equities) whose value is more interest-rate sensitive than short-term assets. Third, consumption by younger households may be less interest-rate sensitive than older (and retired) households because the young are able to adjust labor supply, thereby substituting leisure for consumption in response to monetary policy shocks. In short, younger households tend to pay for their consumption with labor income; whereas, older households rely on assets whose value are interest-rate sensitive. Fourth, older households discount the future more heavily on account of a higher probability of death. This, combined with shorter planning horizons, makes monetary policy shocks feel more permanent for the older households, and induces additional interest-rate sensitivity into their consumption.

To explore the potential role of these life-cycle related effects, we employ household income as a proxy for wealth. We then estimate the VARS and local projections for consumption on households classified by income and age. The highest consumption responses to monetary policy shocks are observed in the oldest and high income groups. We continue our argument that heterogeneity in age and wealth can drive these consumption response patterns with a life-cycle model of consumption, saving, and labor-supply decisions. Preferences in the model are given by Epstein and Zin (1989)–Weil (1989) recursive utility. Finitely-lived people work and earn labor income from ages 25 to 64. From age 65 to (at most) 86, they live on pension income and accumulated assets. Both labor and retirement income are subject to idiosyncratic uncertainty, as is the time of death, which gives people both a precautionary and a retirement, or life-cycle, motive to save. People cannot die with negative net worth, but can borrow or lend during their working years by taking short or long positions in a long-term asset. We adopt a long-term asset to be consistent with actual household net wealth patterns, which are weighted toward long-term assets. Consumption impulse responses to interest rate shocks in the model qualitatively match the age-related pattern of responses in the data–notably, older households have the largest consumption responses.

Our paper contributes to the growing interest in researching the macroeconomic implications of agent heterogeneity. Studies of monetary policy transmission with heterogeneity arising from incomplete markets include, Gornemann et al. (2012), McKay et al. (2016), and Luetticke (2016). Our paper is also related to Auclert (2016) who stresses heterogeneity in the duration of an agent's net worth, and Fujiwara and Teranishi (2007) who embed life-cycle behavior in a New Keynesian model. Di Maggio et al. (2014) study how households of different income levels respond to reductions in mortgage interest payments; whereas, Wong (2016) studies the consumption implications of mortgage debt refinancing by younger households in response to interest rate reductions following negative monetary policy shocks.

Additionally, Bunn et al. (2018) study how monetary policy impacts inequality in the United Kingdom, while Coibion et al. (2017) studies it for the United States. Bielecki et al.

(2018) consider how demographic change affects the interest rate within a New-Keynesian model of monetary policy. A few recent papers have noted a role for wealth effects in explaining consumption volatility. Krueger and Perri (2006) make this point by examining Italian and U.S. data, while Glover et al. (2017) note how changes in asset prices during the last recession disproportionately impacted older households. Storesletten et al. (2007) also consider idiosyncratic shocks in a life-cycle model, but their focus is on explaining the equity premium puzzle.

The paper is organized as follows. The next section describes the data on consumption and monetary policy shocks used in the empirical analysis. Section 3 reports the main empirical results from the structural vector autoregressions and the local projections. Section 4 considers the evidence for the wealth effects. Section 5 presents the life-cycle model and its analysis, and Section 6 concludes.

2 The Data

This section describes the data used in our main empirical analysis, in which we estimate the responses of consumption by age group to unanticipated monetary policy shocks. The consumption data is described in Section 2.1 and the alternative monetary policy shocks are discussed in Section 2.2.

2.1 Consumption Expenditures

The household consumption expenditure data comes from the interview samples of the Consumer Expenditure Survey (CEX) spanning from 1984Q1 through 2007Q4.[1] We collect quarterly household consumption expenditures for 19 broad categories, which we convert into real terms by deflating the expenditure categories by the corresponding price indices from the Consumer Price Index (CPI). We follow Krueger and Perri (2006) in matching categories between the CEX and CPI. Aggregating over the 19 real expenditure components within each household gives our measure of total real consumption by household. The age of the household head is used to classify households into young (25-34), middle (35-64),

[1] U.S. Department of Labor, Bureau of Labor Statistics, Consumer Expenditure Survey, Interview Survey. CEX availability begins in 1984. We end our empirical analysis in 2007 due to the ending of conventional monetary policy in the U.S. resulting from the global financial crisis.

and old (65+) age groups. Our rationale for these age categories is as follows. The 25-34 age group encompasses most first-time home buyers (i.e., the age when long-term asset accumulation begins). People 65 and older are generally retired.

We construct real per capita consumption by dividing total real household consumption by the number of household members. Appendix A gives a detailed description of the construction of the consumption data.

2.2 Monetary Policy Shocks

We consider three alternative measures of identified monetary policy shocks. These shocks are not our own, but were constructed by other researchers. Using different methodologies, the creators sought to identify the portion of changes to the federal funds rate that are both unanticipated and exogenous to current economic conditions. The original monetary policy shock series are monthly. To match the sampling frequency of our consumption data, we aggregate their monthly observations to a quarterly frequency.

The first measure that we consider comes from Barakchian and Crowe (2013) who employ a high frequency identification (HFI) method. Their signal of the policy stance are federal funds futures contracts for the current month and at 1 through 5 months ahead. The information of these 6 horizons is summarized by a factor representation, and the shock is the change in the first factor on the day of an announcement following a Federal Open Market Committee (FOMC) meeting. The changes in the factor is intended to capture the unexpected change in the term-structure of federal funds futures prices induced by policy surprises. Their series begins in 1988Q4 with the establishment of federal funds futures at the Chicago Board of Trade. We refer to this shock series as HFI-TS.[2]

The second monetary policy shock series we use is the instrument employed in Gertler and Karadi (2015) (this is what they call FF4). They employ the HFI approach of Gürkaynak et al. (2005) by using the change in the three-month ahead federal funds futures price within a 30 minute window around an FOMC announcement. The idea with the short window is that the futures price is responding only to FOMC announcements and not to other news. We refer to this shock series as HFI-3.[3]

[2]The TS part reflects the term-structure aspect of the shock.
[3]The HFI-3 and HFI-TS shock series are from Valerie Ramey:
http://econweb.ucsd.edu/vramey/research.html#data, accessed in August 2017.

The third measure is constructed following the approach in Romer and Romer (2004), where they first draw on narrative accounts from FOMC meetings to create the *intended* federal funds rate. Then to control for anticipated movements in the federal funds rate, they regress the change in the intended federal funds rate on unemployment, Greenbook estimates of past and future inflation and real output, and revisions in these forecasts. The residual series are argued to be exogenous to current economic conditions and free from anticipatory movements. We use the Romer and Romer (2004) updated series by Wieland and Yang (2016) and we refer to the narrative/Greenbook series as NG. Their series starts in 1969, but we begin with 1984Q1, the beginning of our CEX consumption series.

Table 1: Summary Statistics for Alternative Monetary Policy Shocks

A. Mean and Standard Deviation			
		Mean	St. Dev.
HFI-TS		-0.001	0.125
HFI-3		-0.045	0.110
NG		0.052	0.277
B. Pairwise Relationships			
x	y	Correlation	$R^2 : y = \gamma + \alpha x$
HFI-TS	HFI-3	0.374	0.127
HFI-TS	NG	0.308	0.092
HFI-3	NG	0.325	0.093
C. Autoregressions			
Lag	HFI-TS	HFI-3	NG
1	-0.076	0.161	0.276
	(-1.017)	(1.393)	(2.920)
2	0.102	0.253	0.120
	(1.075)	(3.421)	(1.075)
3	0.179	0.132	0.083
	(1.934)	(1.388)	(0.847)
4	-0.034	-0.068	0.060
	(-0.338)	(-0.823)	(0.601)
R^2	0.043	0.143	0.154
p-val (Wald)	0.365	0.000	0.003

Notes: In Panel C, t-ratios are in parentheses and are computed by Newey-West. The Wald test is for joint significance of the 4-lag coefficients. The starting dates for the series are as follows: HFI-TS in 1988Q4; HFI-3 in 1991Q1; NG in 1984Q1. The ending dates are 2007Q4.

Table 1 reports basic features of the three shock series through 2007Q4 along with their relationship with each other. The shock means are insignificantly different from zero (Panel A). The NG shocks are about twice the size of HFI-3 and HFI-TS shocks as seen from the standard deviations (Panel A). The information content of the three shock series is different, as the pair-wise correlations hover between 0.3 and 0.4 (Panel B). Truly exogenous monetary policy shocks should be serially uncorrelated. To check this, Panel C of the table shows fitted fourth-order autoregressions to each of the shocks. The HFI-TS shocks come closest to satisfying this criteria, as the Wald test for joint significance of lagged coefficients is insignificant (while it is significant for HFI-3 and NG). Although the autoregressions display some evidence against exogeneity for these latter two shock series, we will impose the assumption of exogeneity in the empirical work.

To facilitate comparisons of the consumption responses across the alternative monetary policy shock series, we normalize each shock to have the same standard deviation as quarterly changes in the real federal funds rate (0.88 percent per annum). Hence, the response to a one standard deviation innovation in the monetary policy shock series is comparable to a one standard deviation unanticipated change in the policy rate.

3 Empirical Results

This section presents our empirical methodology and reports the estimation results. In Subsection 3.1, we employ a standard structural vector autoregression (VAR) approach. In Subsection 3.2, we use local projections as an alternative method. Subsection 3.3 reports the structural VAR and local projection results for non-durable consumption instead of total consumption (durable and non-durable). We report the results of a number of additional robustness checks in Appendix B. Across the three shocks, two estimation methods, and many robustness checks, the weight of the evidence finds that the highest consumption response to monetary policy shocks is by older households. Additionally, the heterogeneity across age groups is quantitatively large.

3.1 Structural Vector Autoregressions

Our structural vector autoregressions (VARs) are similar to those employed by Anderson et al. (2016) and Ramey (2011), who study consumption responses to fiscal policy shocks. The variables in the VAR are the quarterly growth rate of average real per capita consumption, $g_{c,j,t}$ for age group $j = \{\text{young, middle, old}\}$ at time t, the monetary policy shock, s_t, and the real federal funds rate, r_t.[4] Seasonal dummies are included in all regressions, except where noted. We estimate separate VARs for each of the three age groups, young (24-34), middle (35-64), and old (65+). To avoid clutter in this exposition, we suppress the age group j subscript.

The ordering of the variables are consumption growth, the monetary policy shock, and the real federal funds rate. Let $Y_t = [g_{c,t}, s_t, r_t]'$. Suppressing the constants, the VAR is

$$AY_t = \sum_{p=1}^{k} B_p Y_{t-p} + u_t, \tag{1}$$

where the structural error terms are serially uncorrelated with a diagonal covariance matrix, $E(u_t u_t') = D$. Multiplying both sides of equation (1) by A^{-1} gives the reduced form VAR,

$$Y_t = \sum_{p=1}^{k} C_p Y_{t-p} + \epsilon_t, \tag{2}$$

where $C_p = A^{-1} B_p$. The reduced form error vector $\epsilon_t = A^{-1} u_t$ has covariance matrix $E(\epsilon_t \epsilon_t') = \Sigma$. Identification of the structure is obtained by the Choleski decomposition, $A^{-1} \Sigma A^{-1'} = D$ where A is a lower triangular matrix. The ordering of the variables are such that consumption growth responds to monetary policy shocks only with a lag.

We impose additional structure on the empirical model as follows. If the monetary policy shocks are truly exogenous, neither lags of the shock nor lags of other variables should appear in the equation for s_t. We impose this condition with the following zero restrictions. In the matrix A, set $a_{1,2} = 0$ (the $1 - 2th$ element of A), and let B_p be as follows

$$B_p = \begin{pmatrix} b_{p,11} & b_{p,12} & b_{p,13} \\ 0 & 0 & 0 \\ b_{p,31} & b_{p,32} & b_{p,33} \end{pmatrix}. \tag{3}$$

Due to the relatively short time-span of the data, imposing these theoretical restrictions lightens the parameterization of the VAR and preserves degrees of freedom.

[4] The nominal federal funds rate is deflated by the Personal Consumption Expenditure (PCE) price index to obtain the real federal funds rate. The nominal federal funds rate and the PCE come from the Federal Reserve Bank of St. Louis FRED database.

Figure 2: Structural VAR – Cumulated Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

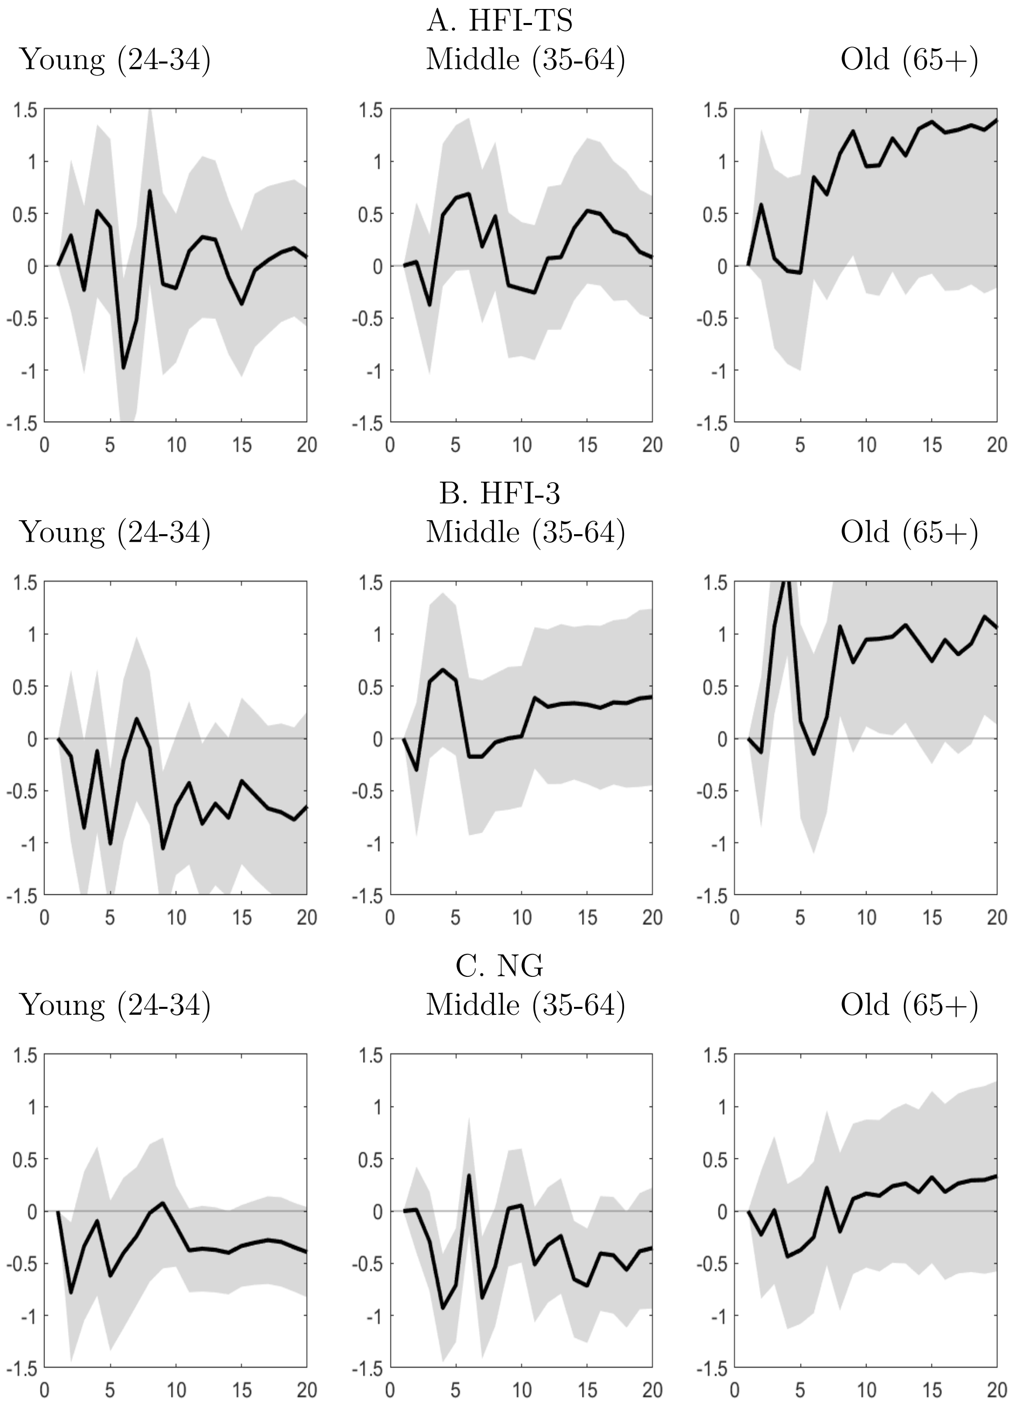

Notes: The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error asymptotic confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

We estimate 8-th ordered VARs. Figure 2 shows impulse response functions (IRFs) of cumulated consumption growth by age group to a negative (expansionary) one standard deviation monetary policy shock. The horizontal axis measures time, in quarters, up to five years after the shock. The vertical axis measures the strength of the consumption response in percent. Responses to the HFI-TS shock are shown in Panel A, responses to the HFI-3 shock in Panel B, and responses to the NG shock in Panel C. The shaded areas are plus and minus one asymptotic standard error confidence bands, commonly used in monetary policy VARs (e.g., Romer and Romer (2004)).

For each of the three monetary policy shocks (Panels A, B, and C), there is striking heterogeneity in consumption responses across age groups. This represents our main empirical finding. Different age groups respond differently to monetary policy shocks, with the old having the largest positive response.

Following an expansionary HFI-TS shock (Panel A), the largest consumption response is by the old. The young and middle age groups have muted consumption responses. After about 8 quarters, consumption for the old households has increased dramatically, and the effect seems to be permanent. The peak consumption response for old households is twice as high as the peak for the middle-aged. As time passes, the impact on young households dissipates.

The HFI-3 monetary policy shock (Panel B) induces the most striking contrast in consumption responses across age groups. The response for the young is generally negative. Consumption of middle-aged households increase two to five quarters after the shock, but their response is short-lived. Consumption for old households increases significantly and again appears to be permanently impacted by the monetary policy shock.

The consumption responses to an expansionary NG shock (Panel C) are relatively subdued in comparison to the other two monetary policy shocks. Young and middle consumption display similarity in timing and magnitudes, both exhibiting modest decreases. Old consumption also initially declines in response to the negative NG shock, but then the response gradually turns positive.

To summarize, consumption responses for the old to each of the three alternative monetary policy shocks are overall largest and long-lasting. Consumption responses of young and middle-aged households are closer to zero. The most striking response heterogeneity across age groups is generated in response to HFI-3 shocks. In all cases the old respond more positively. Monetary policy has distributional effects as the differences across age groups are quantitatively large.

Changing Monetary Policy Effectiveness over Time. The differences in consumption responses across age groups suggest that the evolving demographic composition may increase the effectiveness of monetary policy. To get a sense of the potential impact, we can aggregate our age-specific impulse response functions and combine them with alternative demographic profiles to obtain estimates of policy effectiveness at different points in time: 1990, 2010, and 2030. Let c_A denote aggregate real per capita consumption. The approximate change in $\ln(c_A)$ is

$$N_A \Delta \ln\left(c_{A,t}\right) = N_y \frac{\bar{c}_y}{\bar{c}_A} \Delta \ln\left(c_{y,t}\right) + N_m \frac{\bar{c}_m}{\bar{c}_A} \Delta \ln\left(c_{m,t}\right) + N_o \frac{\bar{c}_o}{\bar{c}_A} \Delta \ln\left(c_{o,t}\right), \tag{4}$$

where estimates of the change in young, middle, and old household real per capita consumption, $\Delta \ln\left(c_{y,t}\right)$, $\Delta \ln\left(c_{m,t}\right)$, and $\Delta \ln\left(c_{o,t}\right)$, come from the estimated VARs, and the number of young (N_y), middle (N_m), old (N_o), and aggregate ($N_A = N_y + N_m + N_o$) population are calculated from the UN World Population Prospects 2017 Revision data (as in Figure 1). For the weights $\bar{c}_y/\bar{c}_A$, etc., we use average age-group consumption shares of aggregate consumption. We abstract from changes in family size and composition within age groups. The exercise here holds the responses and the relative consumption of each age group fixed, but varies the age-distribution (N_y, N_m, and N_o) over time to isolate how changes in the demographic composition impacts monetary policy effectiveness.

Figure 3: Structural VAR – Monetary Policy Effectiveness

A. Difference in the Aggregate Response in 2010 and 2030 Compared to 1990

HFI-TS HFI-3 NG

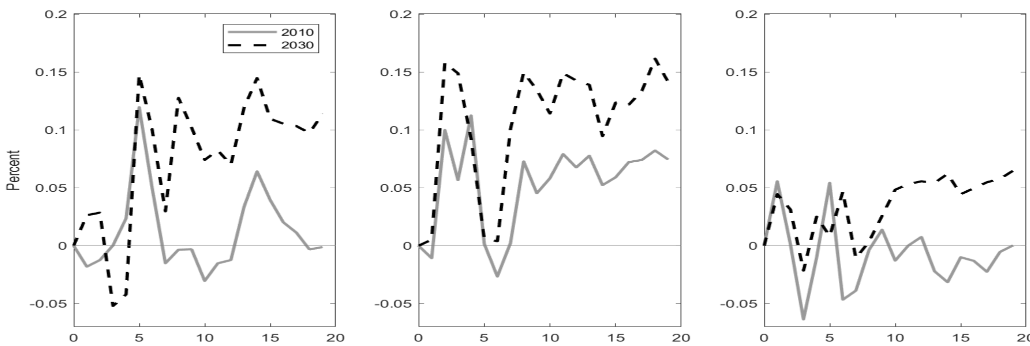

B. Cumulated 20 Quarter Difference in Aggregate Response
between 1990 and 2010 and between 1990 and 2030

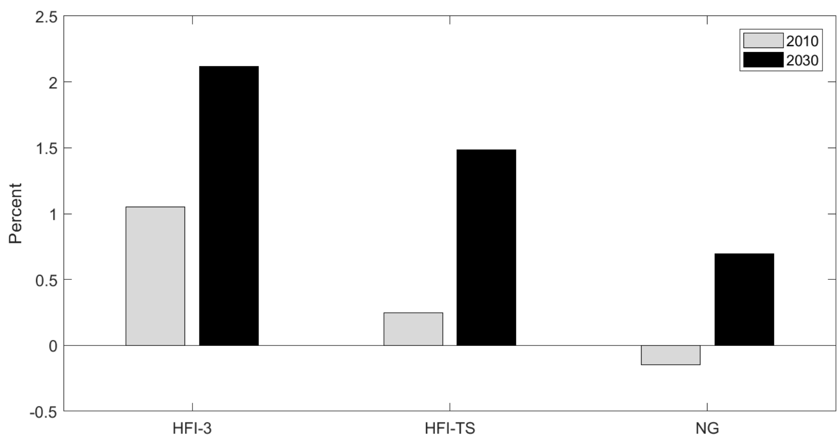

Notes: The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent. We estimate the consumption responses by age group to a monetary policy shock and hold it fixed and then change demographics according to UN World Population Prospects 2017 Revision data. Panel A: Absolute difference between responses in 2030 and 2010 relative to 1990. Panel B: Cumulated percent difference in total consumption in 2030 and 2010 relative to 1990.

Figure 3 displays the results. We estimate the response of cumulated aggregate consumption growth to each of the three monetary policy shocks using population weights in years 1990, 2010, and 2030. Panel A plots the difference in the responses to expansionary monetary policy shocks between 2010 and 1990 and between 2030 and 1990. According to our results, population aging from 1990 through 2030 increases monetary policy effectiveness.

Panel B aggregates the flow consumption differences between 2010 and 1990 estimates and between the 2030 and 1990 estimates over the five years after the shock.[5] These figures highlight how dramatically the age-distribution changes the likely effectiveness of monetary policy, especially for the HFI-TS and HFI-3 shocks. The predicted demographic change is estimated to generate as much as an additional 2.0 percent cumulated change in consumption for a one standard deviation expansionary monetary policy shock in 2030 relative to 1990. While we estimate that population aging from 1990 through 2030 increases monetary policy effectiveness when measured by HFI-TS and HFI-3 shocks, the differential consumption responses to the NG shock is much smaller.[6]

3.2 Local Projections

The structural VARs show a high degree of heterogeneity in consumption responses by households of different ages to monetary policy shocks. To examine the robustness of these results with respect to the empirical procedures, this section uses local projections (Jordà (2005)) as an alternative method. As in the VARs, we include lags of consumption growth, the real federal funds rate, and seasonal dummies in the regressions. The local projections are the sequence of regressions at horizons $h = 1, \ldots, 20$, estimated separately for each age group,

$$\ln\left(\frac{c_{t+h}}{c_t}\right) = \beta_h s_t + a_{h,1} \ln\left(\frac{c_t}{c_{t-3}}\right) + a_{h,2} \ln\left(\frac{c_{t-4}}{c_{t-7}}\right) + \sum_{j=0}^{7} b_{h,j} r_{t-j} + u_{t+h} \tag{5}$$

where c_t is average per capita consumption within an age group at time t, s_t is the identified monetary policy shock, and r_t is the real federal funds rate (Again, we suppress the notation that differentiates across age groups). The coefficient of interest is β_h, which gives the percent change consumption response from time t to $t + h$ from the monetary policy shock at time t. Standard errors are computed by Newey and West (1987). Figure 4 displays the results.

[5] We are calculating the net area between the curves in Panel A and the zero line.

[6] An unappealing feature of the NG shock is that after an initial increase, aggregate consumption declines by 0.14 percent after 20 quarters.

Figure 4: Local Projection – Cumulated Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

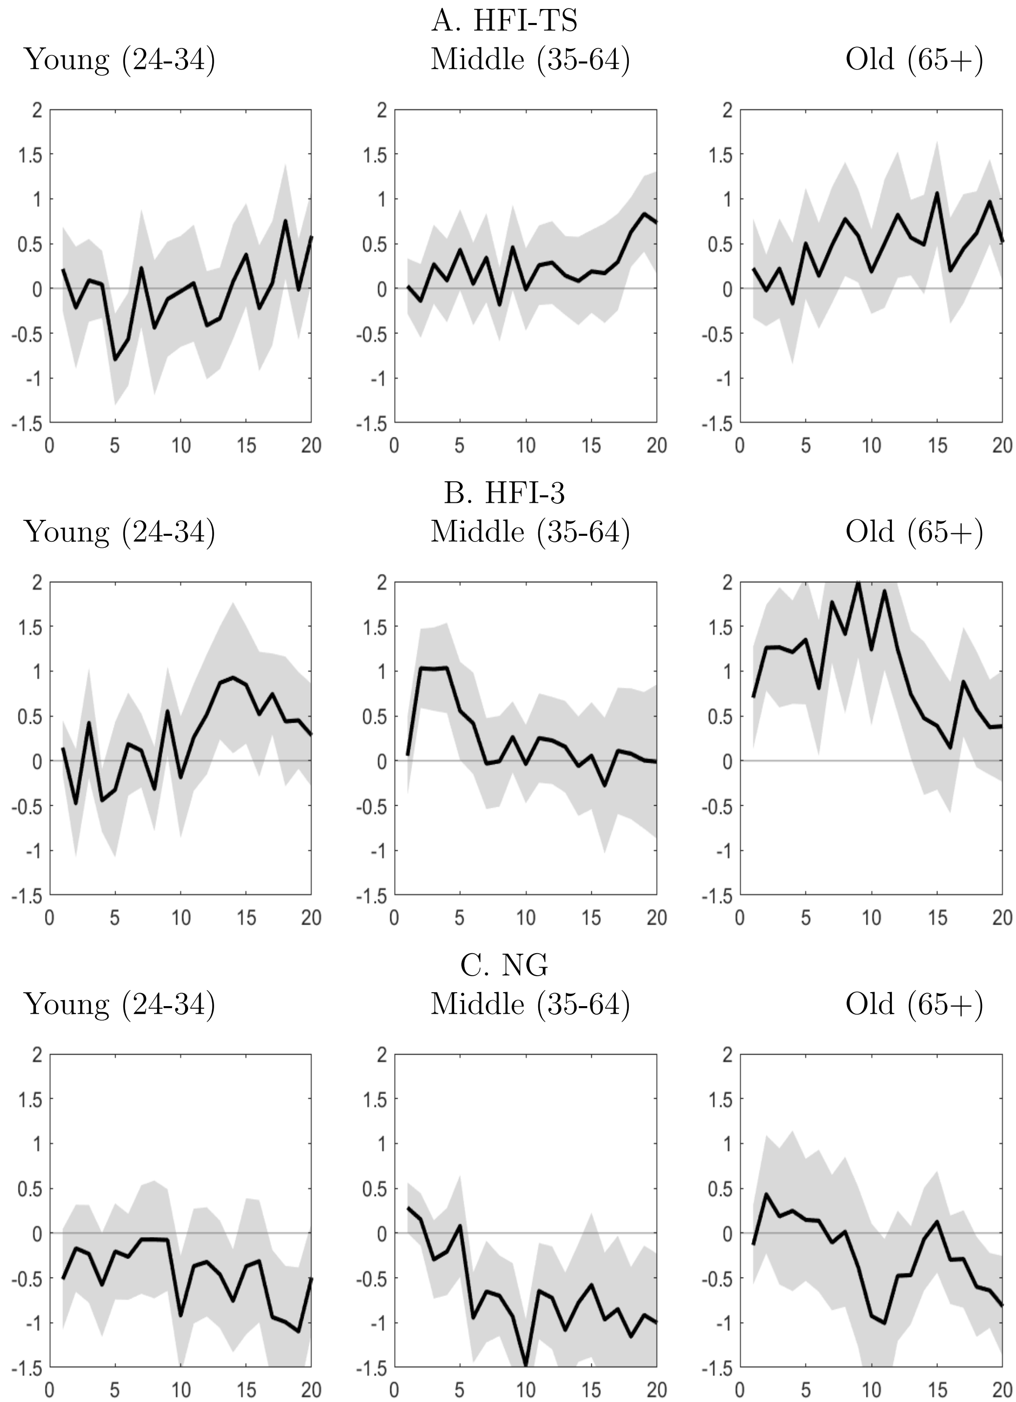

Notes: Estimates are from equation (5). The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error Newey and West (1987) confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

The consumption responses again are seen to vary by age group. The response among the old to the HFI-TS shock is higher than for the younger age groups over much of the following 20 quarters. The old also exhibit the highest response to the HFI-3 shock from quarters 1 through 12, after which the response by the young and old are similar in magnitude. Middle consumption exhibits an initial positive response but it quickly dies out. Looking at panel C, young and middle consumption generally respond negatively to the NG shock. Old consumption increases modestly for about two years after the shock then turns negative.

To summarize the local projections results, old consumption responds strongly positively to the HFI-TS and HFI-3 monetary policy shocks and within the first year after the shock, responds positively to the NG shock. Old consumption generally increases by more than the consumption of middle and young households. The relative ranking of consumption responses between young and middle households is less definitive.

3.3 Response of Non-Durable Consumption

The consumption data studied to this point includes durables, whose purchases may be financed. Researchers and policy makers may also be interested in understanding patterns of non-durable consumption. To examine this, and to verify that our results are not driven entirely by durable expenditures, this subsection examines the non-durable consumption responses across age groups to monetary policy shocks. Our measure of what constitutes non-durable expenditures follows Krueger and Perri (2006).[7]

Figure 5 shows cumulated non-durable consumption growth responses from the structural VAR. As can be seen, both the response patterns and the magnitudes are similar to the total consumption responses displayed in Figure 2. The size ordering of responses across all three shocks is old $>$ middle $\simeq$ young.

Figure 6 shows the corresponding local projection results. Again, the responses to all three shocks are overall similar for non-durables. With the possible exception of the NG shock, the response among the old remains the greatest.

[7] Real non-durable consumption expenditures is the sum of consumption components 1-13 given in the Appendix A Table A–1.

Figure 5: Structural VAR – Cumulated Non-Durable Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

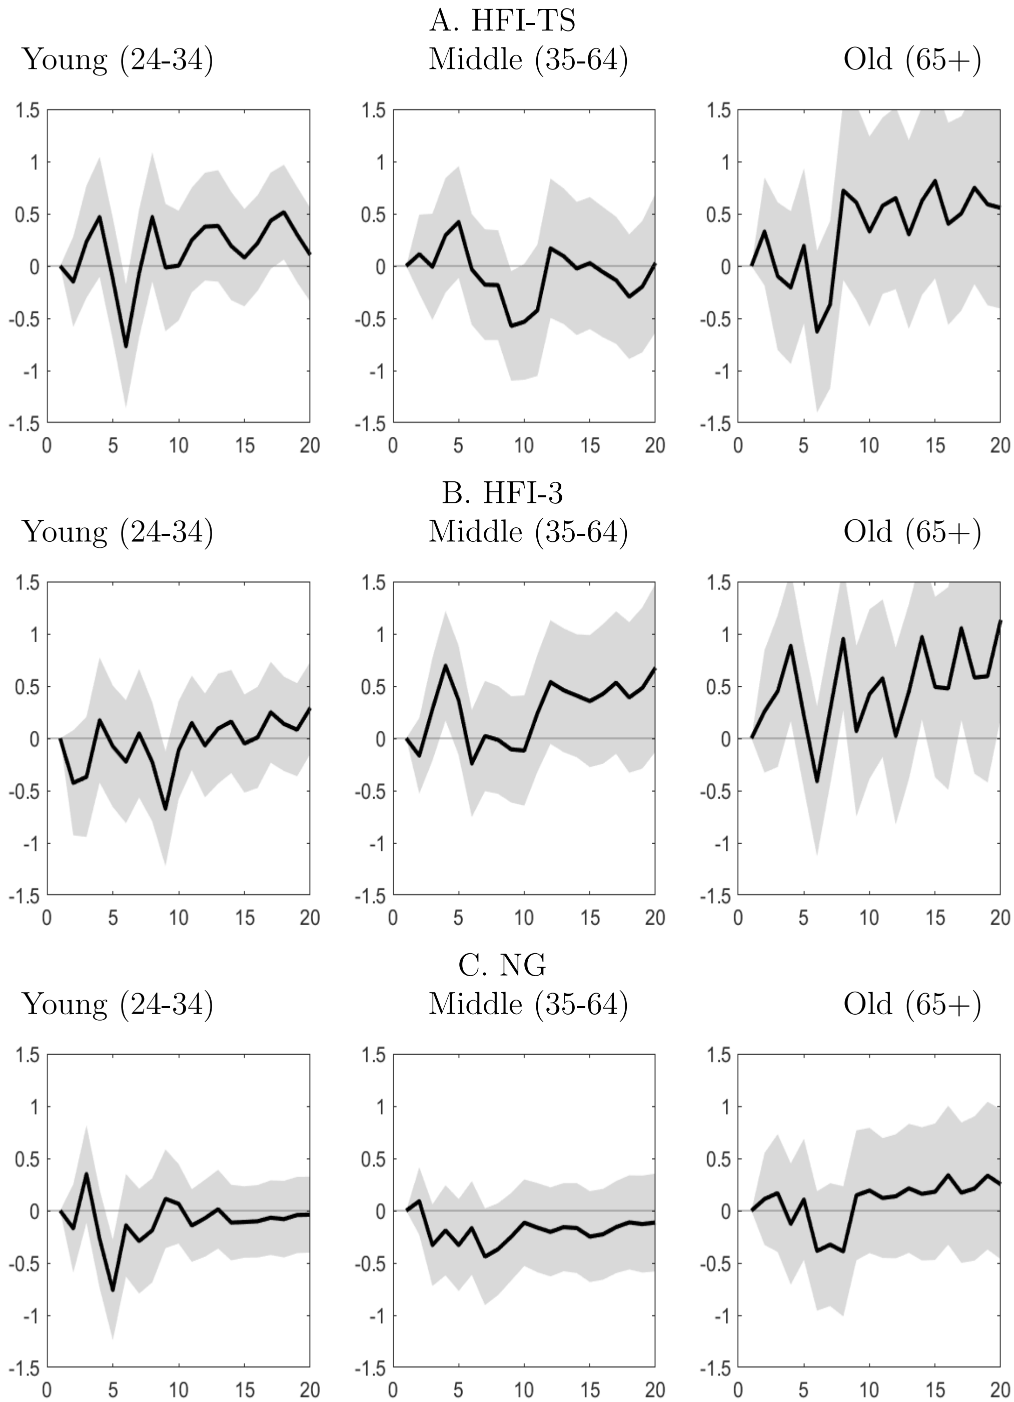

Notes: The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error asymptotic confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

Figure 6: Local Projection – Cumulated Non-Durable Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

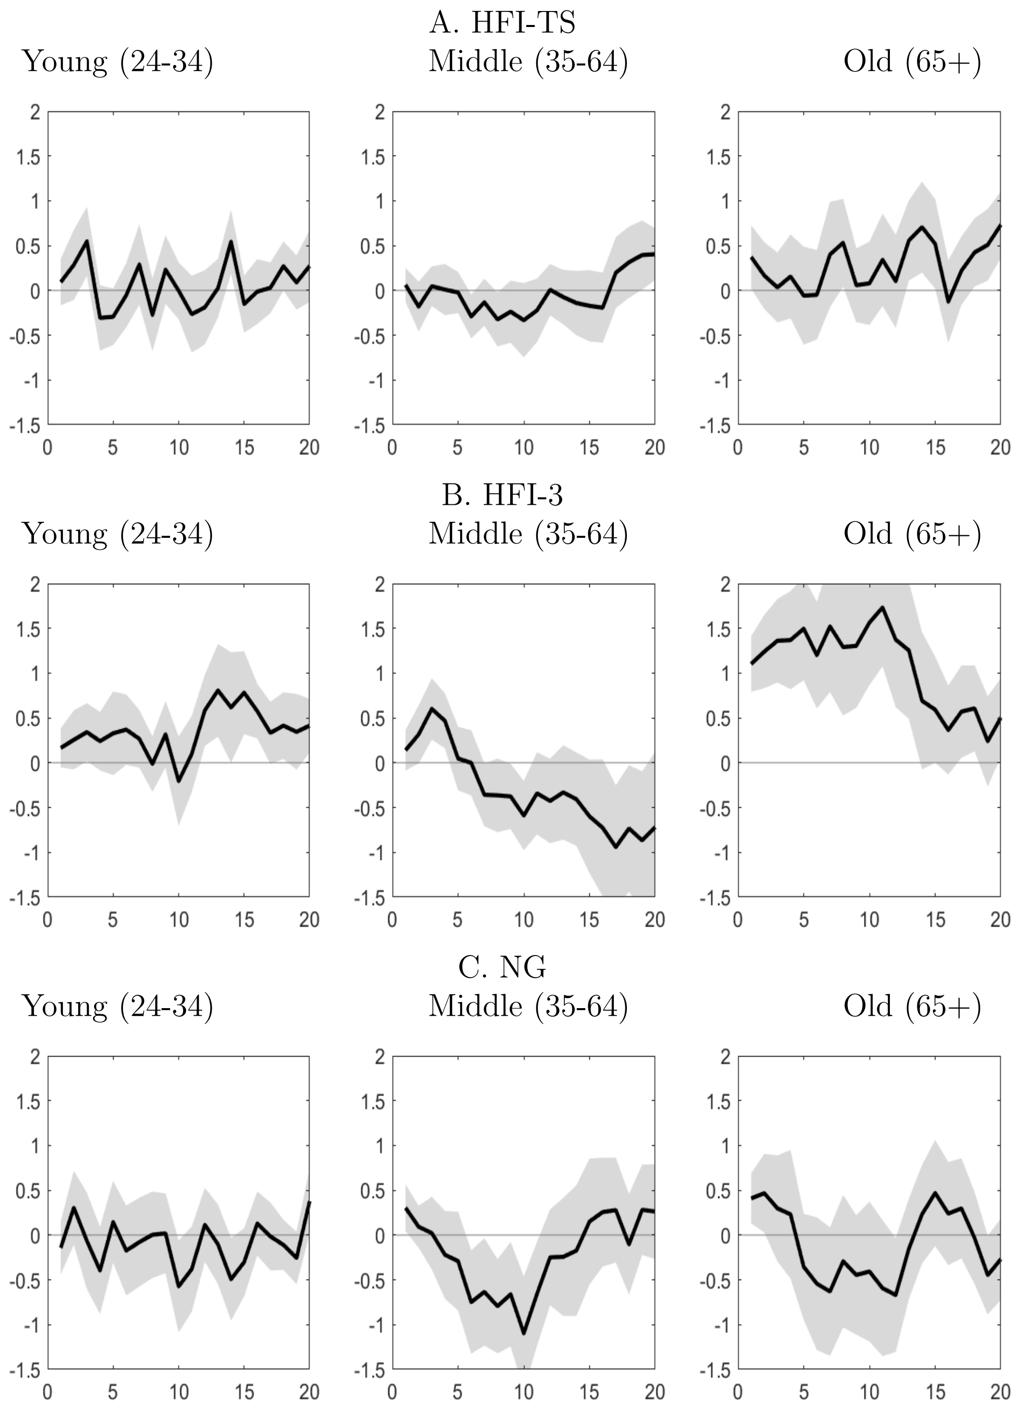

Notes: Estimates are from equation (5). The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error Newey and West (1987) confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

To summarize, the qualitative and quantitative responses of non-durable consumption by age group from the structural VARs and the local projections remain similar to the total consumption responses. For both methods, the general finding is that the old exhibit the largest responses to monetary policy shocks. Based on these results, the differences across age groups do not appear to be driven primarily by the types of goods purchased. Instead, the differences likely come from how consumers finance their consumption as they move through the life-cycle.

Appendix B contains a number of other robustness checks including using more age groups, considering durable goods expenditures (with and without housing), alternative aggregation of household consumption, estimating the structural VAR in levels, changing the variables in the structural VAR, and changing the lags in the structural VAR. Throughout the many specifications and across the three shocks, the message remains that consumption of the old households is most responsive to monetary policy shocks.

In the next section, we provide evidence indicating that different age groups receive their income from very different sources and that the old hold a greater share of their wealth in long-term assets.

4 Age, Income, Wealth, and Portfolio Composition

The previous section showed that consumption among the old is most responsive to monetary policy shocks. A potential factor contributing to the observed patterns is that retired older households, who typically live off of wealth (rather than labor income), have portfolios whose value are tilted toward more interest-sensitive long-term assets. Hence, a potential mechanism behind the variation in consumption response is wealth heterogeneity across age groups. While old households adjust to the monetary policy shock induced wealth shock primarily by adjusting consumption, the wealth effect is less important for young households because they have lower net wealth. Additionally, younger households have an extra margin to adjust in response to monetary policy shocks through their choice of labor supply.

This section explores the relation among age, income, and wealth using data from the 1989, 1998, and 2007 waves of the Survey of Consumer Finances (SCF). The span of the sample approximately overlaps the time coverage of our CEX sample.[8]

Labor Income and Household Age. To see where in the life-cycle labor income is replaced by other sources, Figure 7 shows the median and mean wage income as a share of total income by 5 year age groups. Both the median and mean shares begin a rapid decline around age 55. By age 65, the median share of labor income is zero while the mean share lies in the 20 to 30 percent range, depending on the survey year. Retirement, whether voluntary or involuntary, takes place for most people before age 65. The typical older household does not receive much labor income and must pay for consumption from other sources.

Figure 7: Median and Mean Wage Income as a Share of Total Income by Age: 1989, 1998, and 2007

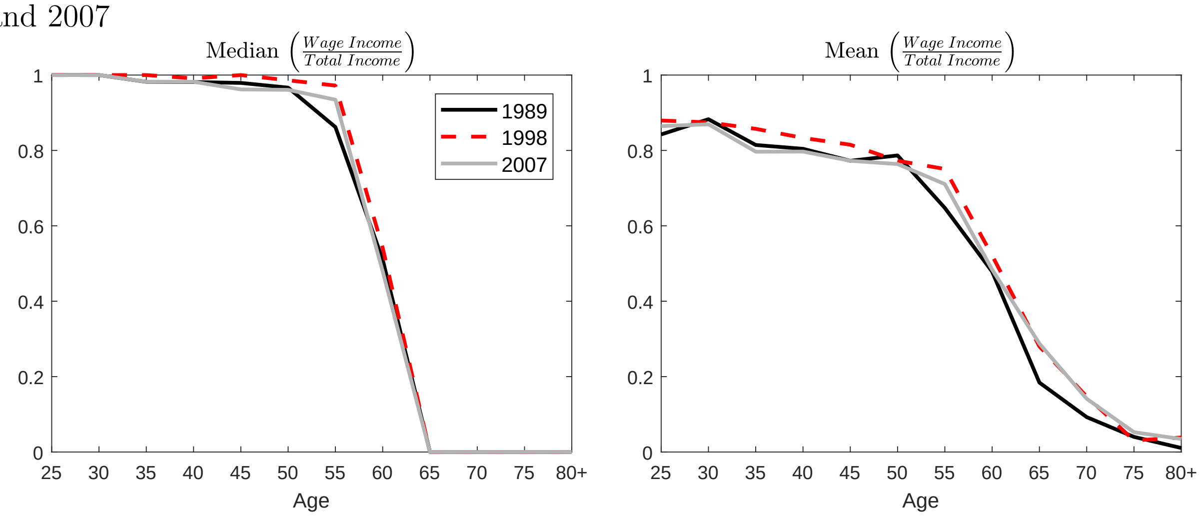

Notes: Median and mean wage income as shares of total income are given by 5 year age group and year. Data is from the Survey of Consumer Finances.

Net Wealth and Portfolio Composition by Household Age. Here, we examine net worth and the composition of long and short-term assets across households of different ages. For each sample year and age group, we construct three measures of net asset positions. *Net Worth* is total assets (as stated by the SCF) minus total debts. Table 2 reports the median net asset position within each age group. Since older age groups hold more wealth, they have more exposure to valuation changes induced by monetary policy shocks.

We also construct an alternative net wealth measure, *Net Long Term 1.* This is the sum of net property equity (value of properties, including own residence, less outstanding

[8]The 1989 SCF survey is the first that allows us to identify holdings in stock mutual funds and annuities.

Table 2: Median Net Assets by Age Relative to Median Net Assets

Age Group	1989	1998	2007
25-34	0.217	0.216	0.169
35-64	1.633	1.281	1.344
65+	1.655	1.927	1.792

Notes: *Net Worth* = total assets - total debts. Figures are median net worth by age and year divided by median net worth of all households in that year. Data is from the Survey of Consumer Finances.

Table 3: Median Net Long Term Assets as a Share of Net Worth

	$\frac{Net\ Long\ Term\ 1}{Net\ Worth}$			$\frac{Net\ Long\ Term\ 2}{Net\ Worth}$		
Age Group	1989	1998	2007	1989	1998	2007
25-34	0.148	0.078	0.130	0.190	0.108	0.139
35-64	0.568	0.437	0.485	0.603	0.471	0.502
65+	0.573	0.625	0.676	0.625	0.670	0.691

Notes: *Net Worth* = total assets - total debts. *Net Long Term 1* = Net property equity (value of primary residence + other residential property - remaining mortgage and debt secured by primary and other residential property + net equity in non-residential real estate) + stocks + stock mutual funds + annuities. *Net Long Term 2* = *Net Long Term 1* + non-stock mutual funds (bond and other mutual funds, not including money market funds) + directly held bonds. Data is from the Survey of Consumer Finances.

debt on the properties), stock holdings, stock mutual funds, and annuities - a measure of long term, interest sensitive assets. Finally, *Net Long Term 2* is more comprehensive and includes the holdings in *Net Long Term 1* plus non-stock mutual funds (not including money market funds) plus directly held bonds of all types.

Table 3 shows the fraction of net worth formed by *Net Long Term 1* and *2* by age group and year. A household with a larger fraction has a higher proportion of net wealth composed of long-term assets, and therefore higher exposure to interest rate fluctuations. We see that the share of long-term assets in net worth increases rather consistently with age. The largest component of net asset holdings for each age group is in property equity. While this is an important component of long-term assets, we also find the same relationship of increasing shares of average interest-sensitive asset holdings by age when property equity is excluded in *Net Long Term 1* and *Net Long Term 2* and Net Worth. We also find that the oldest age groups hold (marginally) more stocks as a share of financial asset holdings than the younger cohorts. Generally speaking, the data show not only that older households have higher net wealth, but also that the value of those assets are more interest-rate sensitive.

Figure 8: Log Net Wealth and Log Income in 1989, 1998, and 2007

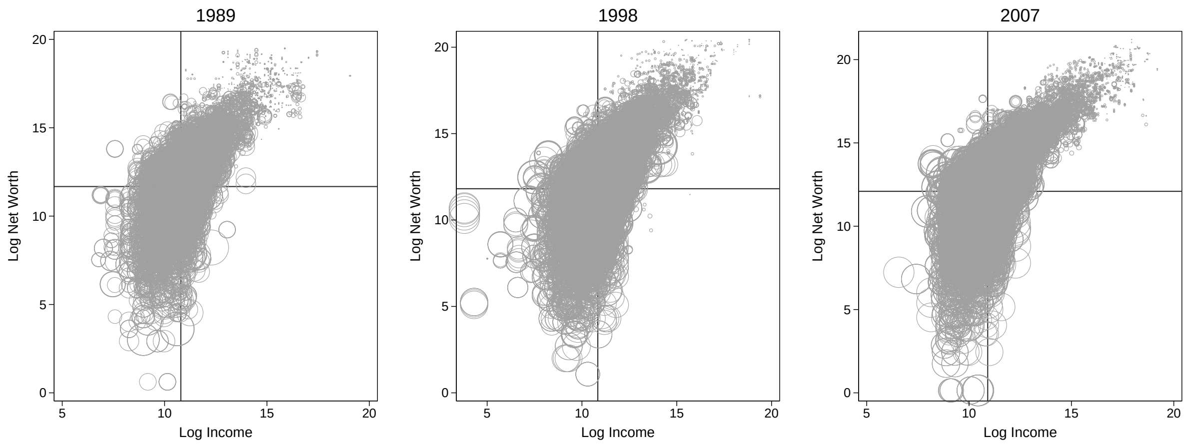

Notes: Data is from the *Survey of Consumer Finances*.

We would like to estimate the VARs using consumption stratified by age and wealth, but the CEX data does not contain a wealth measure suitable for such analysis. As an alternative, we classify households using income as a proxy for wealth. To motivate doing so, Figure 8 plots log net worth against log income from the SCF in 1989, 1998, and 2007.[9] The size of each observation is the representative SCF weights for that household. The vertical and horizontal lines are, by year, the median log income and log net worth. The correlations between the variables are 0.56 in 1989, 0.58 in 1998, and 0.60 in 2007.[10]

The strong positive relationship between income and wealth imparts confidence that household classification by income reasonably proxies for classification by wealth. To implement the investigation, for each age group, we split households into high and low income, where high income are those in the top income decile and low income are those in the bottom decile. The variable of interest is the growth rate of average per capita consumption within each age-income group.

Figure 9 shows the cumulated consumption impulse responses from the structural VAR by age and income to one standard deviation expansionary monetary shocks. For the old, HFI-TS and NG shocks generate much larger positive consumption responses from high

[9]Using logs drops those with negative net worth, but due to the the extreme wealth observations at the tails, we use logs to visualize the relationships.

[10]In line with these estimates, Ríos-Rull and Kuhn (2016) find that the correlation of wealth and income is 0.58 in the 2013 SCF (see Table 19).

income households. The contrast is less stark for the HFI-3 shock. The difference between high income and low income middle households is also less pronounced, and there is almost no difference for younger households. A possible reason for the lack of clarity for young and middle-aged households may be that the wealth shocks feel less important due to their relatively lower net asset holdings and greater labor income.

Figure 10 shows the corresponding local projection results by age and income. Here, the high income old have a higher response to HFI-TS and NG shocks. Results show little difference in responses between high and low income for young and middle-aged households.

The results by income fairly consistently show that for the old, there is a higher consumption response to monetary policy shocks for the high income households. These results from stratifying households by income and age suggest heterogeneity in wealth-effects as a potential driving mechanism. The next section addresses this question with a partial-equilibrium life-cycle model.

Figure 9: Structural VAR – Cumulated Consumption Growth IRF by Age Group and Income to Expansionary Monetary Policy Shock

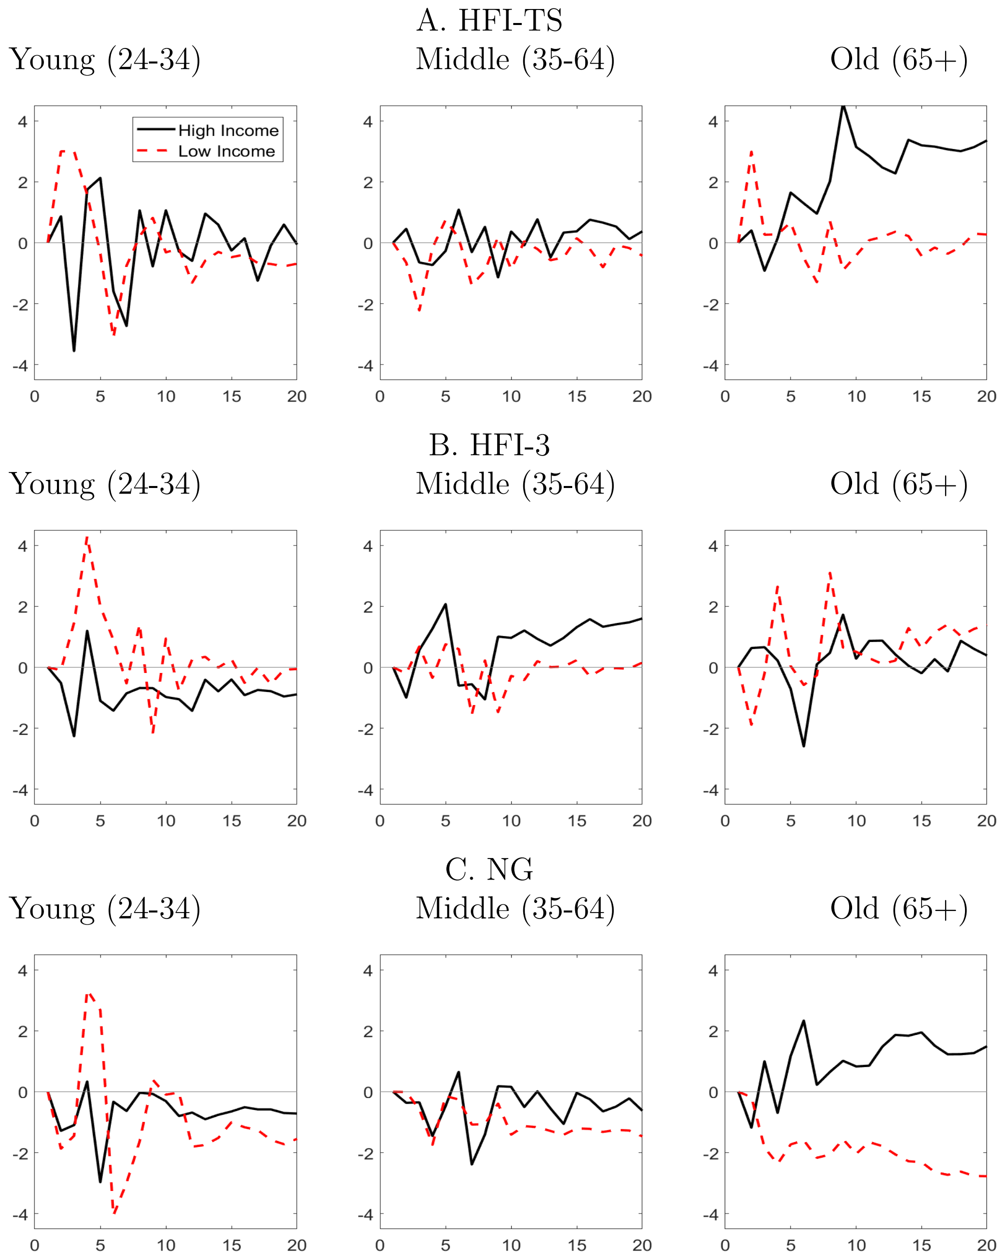

Notes: High (low) income is average per capita consumption in the top (bottom) income decile by age group. The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error asymptotic confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

Figure 10: Local Projections – Cumulated Consumption Growth IRF by Age Group and Income to Expansionary Monetary Policy Shock

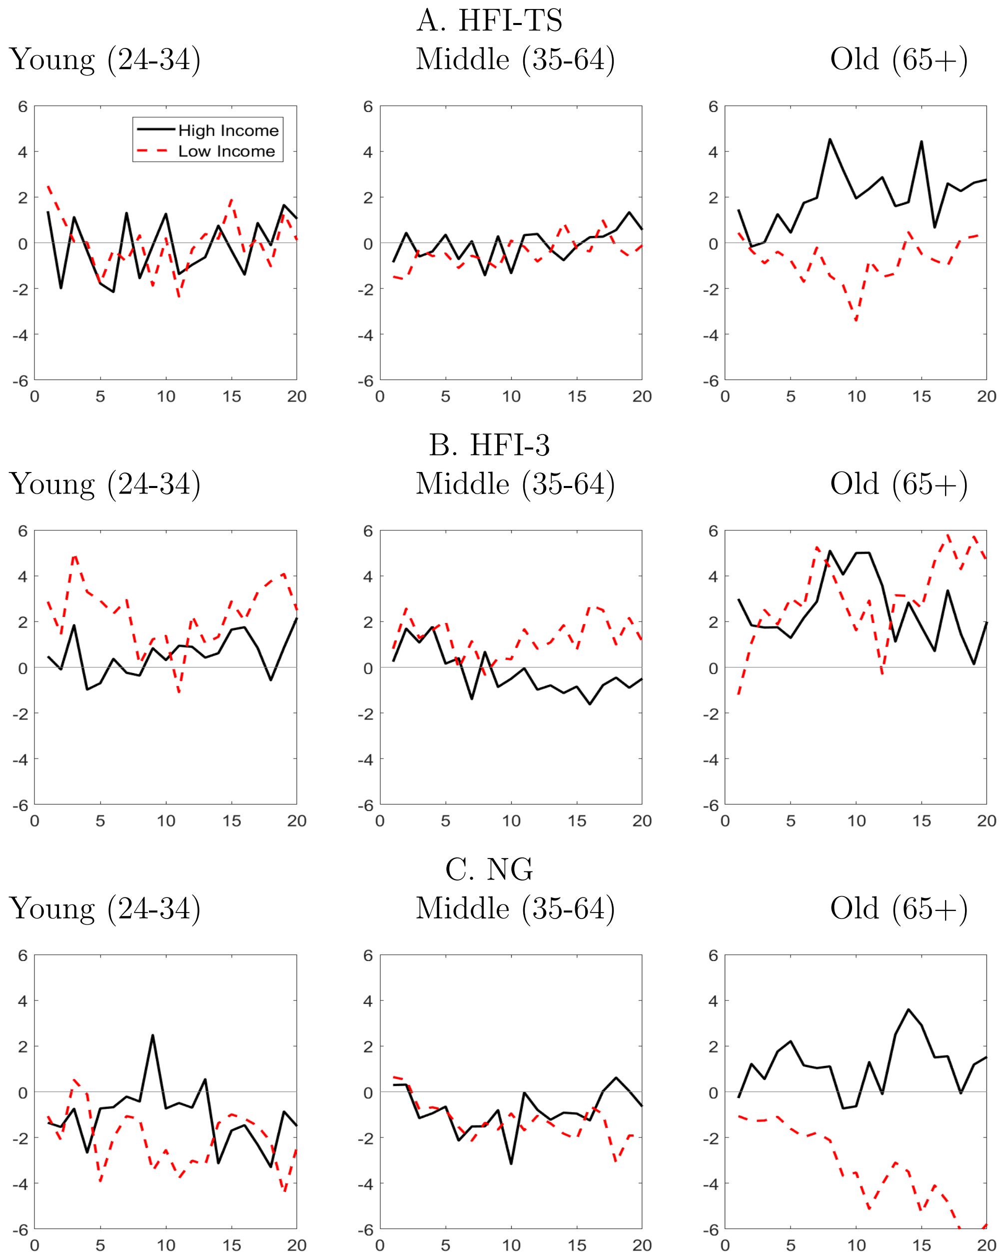

Notes: High (low) income is average per capita consumption in the top (bottom) income decile by age group. The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are ± one standard error Newey and West (1987) confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

5 Wealth-Effects and Consumption Heterogeneity in a Life-Cycle Model

This section illustrates how heterogeneity in wealth-effects, labor-supply choices, planning horizons, and discounting of the future might explain the observed consumption dynamics across age groups in a model of overlapping generations of finitely-lived households. The model agents save by accumulating long-term bonds both for retirement and because they face income shocks, and preferences are given by Epstein and Zin (1989)–Weil (1989) recursive utility. With these key ingredients, the model can replicate the higher consumption response among the old households.

Model agents live a maximum of 86 years (344 quarters or periods in the model). People begin making economic decisions when they are 25 years old, enter economic life with no assets, and there are no bequests. At each point in time, 248 different decision making cohorts are alive at different stages of the life cycle.

We categorize people into the same three age groups as in the empirical section. Young (25-34 years) and middle (35-64 years) aged people receive exogenous, risky labor income (W) and decide their labor supply (L), consumption (C), and net asset positions (A). Working age households can borrow, but households are not allowed to die with debt. When people turn 65, they retire, face uncertain death, and live off of reduced pension income (S) and accumulated assets. To conform with the long-term interest-sensitive assets that dominate household portfolios (as seen in the previous section), the long-term asset in the model is a consol bond. We begin with a description of the exogenous income process.

5.1 The Income Process

We adopt the permanent-transitory income component model employed by Choi et al. (2017), who in turn draw upon Zeldes (1989), Carroll (1992), and Carroll (1997). Let there be N individuals per cohort. In each quarterly time period t, living cohorts are indexed by $z \in [1, 248]$. Cohort $z = 1$ begins economic life as a 25 year old household, cohort $z^* = 161$ are newly retired, and cohort $z = 248$ are in the last quarter of life. Working age household i of cohort $z < z^*$ draws labor income ($W_{i,z,t}$) and each retiree ($z \geq z^*$) draws pension income ($S_{i,z,t}$). Both labor and retirement income have a permanent component ($Y_{i,z,t}$) and a transitory component ($e^{u_{i,z,t}}$). The idea behind subjecting retiree pensions to permanent income risk is to capture events such as bad health shocks that generate large out-of-pocket medical expenses, while recognizing that these are not utility enhancing consumption expenditures. That is, we substitute modeling health shocks with old-age pension shocks.

The transitory income shock $u_{i,z,t}$, is drawn from a mixture of a normal distribution and a low-probability event in which there is zero income for that quarter

$$u_{i,z,t} = \begin{cases} N\left(\mu_u, \sigma_u^2\right) & \text{with probability } \ (1-p) \\ -\infty & \text{with probability } p \end{cases} \tag{6}$$

where p is the probability of drawing zero income, and $\mu_u = \frac{-\sigma_u^2}{2} - \ln\left(1-p\right)$. This mixture of distributions is frequently employed to model the empirical features of income data which is approximately log-normally distributed except for a concentration of observations at the lower tail. Recalling that cohort z^* has just retired, the labor income for people in their working years is

$$W_{i,z,t} = Y_{i,z,t} e^{u_{i,z,t}} \text{ for } z < z^* \tag{7}$$

and pension income for retired cohorts is

$$S_{i,z,t} = \begin{cases} Y_{i,z,t} & z = z^* \\ Y_{i,z,t} e^{u_{i,z,t}} & z > z^* \end{cases}. \tag{8}$$

During the working years, wage growth is driven in part by a common secular component, whose gross growth rate is M_g, and also by the individual's advance up the age-earnings profile, whose gross growth rate for age z is G_z. In retirement, both M_g and G_z become 1. Let $n_{i,z,t} \overset{iid}{\sim} N\left(\mu_n, \sigma_n^2\right)$ be the shock to permanent income $Y_{i,z,t}$. Then, the life-cycle of earnings evolves according to

$$Y_{i,z,t} = \begin{cases} Y_{i,z-1,t-1} M_g G_z e^{n_{i,z,t}} & z < z^* \\ p_{rr} Y_{i,z^*,t-1} & z = z^* \\ Y_{i,z-1,t-1} e^{n_{i,z,t}} & z > z^* \end{cases} \tag{9}$$

where p_{rr} is the replacement rate on pension income. Note that in the retirement period, the household receives p_{rr} with certainty, after which income resumes its risky evolution.

Estimates of the income process. We estimate the income process from biennial waves of the Panel Study of Income Dynamics (PSID) data between 1986 and 2007, which aligns closely to our CEX sample period. We use the same definition of household income in the PSID as Blundell et al. (2008) and Storesletten et al. (2007), and our estimation method follows Choi et al. (2017), who build on Zeldes (1989), Carroll (1992), and Carroll (1997). We first remove the aggregate time trend, predictable life-cycle or occupation dependent fluctuations, and household fixed effects from the income data, then we use the remaining variation to estimate the parameters (σ_n, σ_u, p). The gross secular growth rate of household income M_g is given by average real income growth across households over the entire sample period. We estimate the age-income profile G_z using variation in income by age. The age-income profile is assumed to be constant over time.

Table 4 reports the estimated parameters for the income process. The data allow direct estimation of the age-income profile for household heads aged 25-60. Given these estimates, values for ages 61-64 are 'forecasted' with a cubic trend regression. Income peaks at age 51. Macroeconomic income growth is virtually nil, with an annual growth rate of 60 basis points.

The remaining parameters (p, σ_u, σ_n) are estimated separately for young, middle, and old age groups. There are modest differences across age groups. The old are most likely to experience a near zero income event with $p = 0.31$, whereas volatility of permanent income is highest for the young, $\sigma_n = 0.14$.

Heathcote et al. (2010) (pages 698, 699) obtain estimates of the standard deviations for the transitory and permanent components of wages (not household income) that are very similar to ours. So, the distinction between labor wages and total income may not matter substantially for prime age workers (as is also suggested by the empirical work above).

Table 4: Annual Income Process Estimates

A. Gross Growth from Age-Income Profile						
Age	G_z	Age	G_z	Age	G_z	
24	1.065	38	1.016	52	0.999	
25	1.061	39	1.016	53	0.994	
26	1.053	40	1.016	54	0.988	
27	1.046	41	1.016	55	0.981	
28	1.040	42	1.016	56	0.973	
29	1.034	43	1.016	57	0.964	
30	1.030	44	1.016	58	0.953	
31	1.026	45	1.015	59	0.942	
32	1.023	46	1.014	60	0.929	
33	1.021	47	1.013	61	0.915*	
34	1.019	48	1.012	62	0.899*	
35	1.018	49	1.009	63	0.881*	
36	1.017	50	1.007	64	0.862*	
37	1.016	51	1.003			
B. Gross Secular Growth						
M_g	1.006					
C. Process Parameters						
	Young	Middle	Old			
	25-35	36-64	65+			
p	0.185	0.231	0.308	(%)		
σ_u	0.471	0.467	0.482			
σ_n	0.144	0.120	0.126			

Notes: * are values forecasted by cubic trend. M_g is gross secular income growth, G_z is age-specific income growth, p is the probability of zero income, σ_u is the standard deviation of transitory income, and σ_n is the standard deviation of permanent income.

5.2 Preferences and Budget Constraints

Households have recursive, non-expected utility, following Epstein and Zin (1989) and Weil (1989). Let $C_{i,z,t}$ denote consumption of household i, with age z, at time t. Labor supply is $L_{i,z,t}$ and, normalizing the time endowment to 1, leisure is $(1 - L_{i,z,t})$. Utility of working age households $z < z^* = 161$, is

$$V_{i,z,t} = \left\{(1-\beta)\left(C_{i,z,t}^{\nu}\left(1-L_{i,z,t}\right)^{1-\nu}\right)^{(1-\rho)} + \beta\left[E_t\left(V_{i,z+1,t+1}^{1-\gamma}\right)^{\frac{(1-\rho)}{(1-\gamma)}}\right]\right\}^{\frac{1}{(1-\rho)}} \quad (10)$$

where $0 < \rho$, $0 < \gamma$, $0 \leq \nu \leq 1$, and $0 < \beta < 1$ is the subjective discount factor. The parameter ρ^{-1} is the intertemporal elasticity of substitution. If there were no labor choice, γ would be relative risk aversion. However, allowing labor choice gives households an additional margin along which to respond to wealth shocks. A decline in wealth can partially be absorbed by working more in addition to cutting back on consumption. Hence, with variable labor, Swanson (2018) shows that risk aversion depends on a combination of parameters controlling both consumption and labor margins. For this specification of utility, equation (10), Swanson (2018) shows that relative risk aversion is[11]

$$RRA = \gamma + (1-\gamma)\left(\frac{\gamma-\rho}{1-\rho}\right). \tag{11}$$

Households live in an incomplete markets environment. Neither contingent claims nor insurance instruments are available to allow households to insure against idiosyncratic income risk.[12] The available non-human asset is a non-state contingent long-term (consol) bond, that pays one unit of consumption each period forever. Our intent is for the long-term bond to mimic the interest-rate sensitivity of home equity, which forms a major part of the typical U.S. household's portfolio, without modeling specific frictions (e.g., lumpiness, down payments, mortgage refinance, housing services in utility) that are associated with housing.[13] Additionally, our previous analysis of the SCF data revealed that households (especially older ones) hold many other long-term interest-rate-sensitive financial instruments.

A working-aged household can borrow or lend by going long or short the consol.[14] The net number of bonds held by the household is $A_{i,z,t}$. Upon retirement, households face the possibility of death. Households must have non-negative assets in retirement to ensure that they do not die with debt. A borrower, $A_{i,z,t} < 0$, pays one unit of consumption per bond while a saver receives one unit of consumption per bond and the price of the bond is the inverse of the interest rate, $P_t^a = 1/r_t$. Current wealth for working-aged households consists of the net bond coupon ($A_{i,z,t}$) plus the market value of the bonds plus labor income less

[11] The intertemporal elasticity of substitution over deterministic consumption paths is the same as for expected utility, ρ^{-1}.

[12] Our estimation of income shocks in the data was net of all transfers and thus corresponds to the notion of un-insurable risk in the model.

[13] Wong (2016) models mortgage refinance as a monetary policy transmission channel.

[14] Appendix D reports the results when the saving instrument is instead a one-period bond.

consumption. Their budget constraints are

$$P_t^a A_{i,z+1,t+1} = A_{i,z,t} + P_t^a A_{i,z,t} + L_{i,z,t} W_{i,z,t} - C_{i,z,t}, \tag{12}$$

which can be written in a more familiar form,

$$A_{i,z+1,t+1} = A_{i,z,t} + r_t \left(A_{i,z,t} + L_{i,z,t} W_{i,z,t} - C_{i,z,t} \right). \tag{13}$$

Retired households supply no labor. Utility for these households, aged $z^* \leq z < Z = 248$, is

$$V_{i,z,t} = \left\{ (1 - \beta\delta_{z,t}) \left(C_{i,z,t}^{\nu} \right)^{(1-\rho)} + \beta\delta_{z,t} \left[E_t \left(V_{i,z+1,t+1}^{1-\gamma} \right)^{\frac{1-\rho}{1-\gamma}} \right] \right\}^{\frac{1}{1-\rho}}, \tag{14}$$

where $\delta_{z,t}$ is the cohort z specific probability of surviving to age $z+1$. In the last quarter of life, $z = Z = 248$, utility is

$$V_{i,Z,t} = (1 - \beta\delta_{Z,t}) \left(C_{i,Z,t}^{\nu} \right)^{(1-\rho)}. \tag{15}$$

Retired households face budget constraints

$$A_{i,z+1,t+1} = A_{i,z,t} + r_t \left(A_{i,z,t} + S_{i,z,t} - C_{i,z,t} \right)$$

with $A_{i,z,t} \geq 0$.

5.3 Solution and Parameterization

To solve the model, we discretize the state space and obtain policy functions for the stationary model. The household's problem is solved by working backwards from the last period of life. The implied level (un-normalized) values are then obtained by multiplying by permanent income. Appendix C describes the stationary transformation.

An exogenous short-term interest rate, independent of household income, follows an AR(1) process which we estimate from the data on the real federal funds rate and discretize following Tauchen and Hussey (1991). From the short-term interest rates, we employ the expectations theory of the term structure of interest rates, with two modifications, to obtain the long-term rate. Because the consol rate implied by the expectations theory is constant, our first modification is to approximate the consol rate with the implied 10 year yield. Second, because the expectations theory generates a flat yield curve, we add a term premium of 1.309 percent, which is the average 10 year term premium found in the data from 1990 to 2007.

Retirees receive 40% of the labor income from their last period of work as a pension (the replacement rate is $prr = 0.4$). Baseline utility function parameters are $\beta = 0.9962$, $\rho = 10$, $\gamma = 12$, and $\nu = 0.5$. This gives an annualized rate of time preference of 1.54%, an intertemporal elasticity of substitution of 0.1, and a moderate degree of risk aversion. With infinitely lived agents, a value of $\nu = 0.36$ typically gives a steady state choice of time worked at 1/3 of the time endowment. We set ν at a slightly higher value.

5.4 Model Impulse Responses

We run the economy simulation for 300 periods (quarters). After 248 time periods, the economy is fully populated by the full complement of cohorts. Each cohort consists of 10,000 individuals. The impulse event is a decline in the interest rate. The shock is a 0.16 percent decline in the long-term bond rate, generated by a decline in the short-term rate.

We simulate un-normalized responses of log consumption, log labor supply, asset quantities, and asset values for each individual to a negative interest rate shock.[15] We then take the mean within each age cohort, and then the mean within each of the three broad age groups. Figure 11 shows the *relative* responses of mean log consumption, mean asset holdings, mean asset value, and mean log labor supply across age groups to the expansionary interest rate shock. In this figure, if the shock occurs at time $t*$, the relative responses are $\ln(C_t/C_{t*})$ for consumption, $\ln(L_t/L_{t*})$ for labor, A_t/A_{t*} for asset holdings, and $(P_t^a A_t)/(P_{t*}^a A_{t*})$ for asset value. Panel A shows that old consumption is the most responsive to the negative interest rate shock. The relative consumption responses for middle and young are smaller and similar to each other. Our life-cycle model can replicate the main qualitative features of the consumption responses estimated from the CEX data.

[15] Assets are not logged since young and middle households can borrow, resulting in negative values for assets.

Figure 11: Relative Responses to Negative Interest Rate Shock

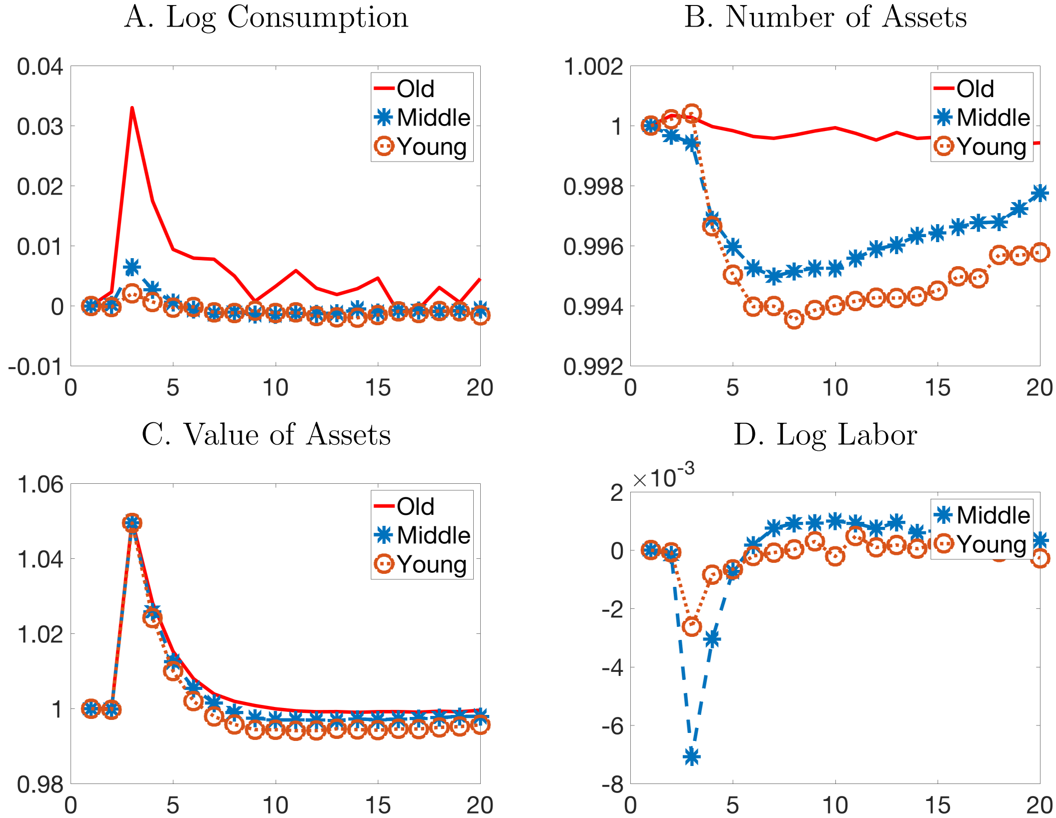

Notes: The figure shows the simulated relative responses by age group to a decline in the long-term bond rate. If the shock occurs at time $t*$, the relative responses are $\ln(C_t/C_{t*})$ for consumption, $\ln(L_t/L_{t*})$ for labor, A_t/A_{t*} for asset holdings, and $(P_t^a A_t)/(P_{t*}^a A_{t*})$ for asset values. The horizontal axis indicates the number of quarters for up to five years after the shock.

In the model, the mechanism works through the wealth effect and the labor-supply margin. Even though the middle and young draw down relatively more assets than the old (Figure 11, Panel B), the old hold far more assets. Because of this, while the relative response patterns of asset values (Figure 11, Panel C) held by the different age groups are roughly the same, the old get a much bigger increase in wealth. Finally, the relative labor responses of both middle and young households are very similar (Figure 11, Panel D). The interest rate decline causes both consumption and leisure for working-age households to increase.

Figure 12 shows the histograms of the asset holding positions for the three age groups. The distributions for middle and old households are heavy in the right tail (note the difference in scale). From the figure, we also see that the consumption response ordering follows from the old having the highest net worth, followed by middle, then young households.

Figure 12: Net Asset Position Histograms

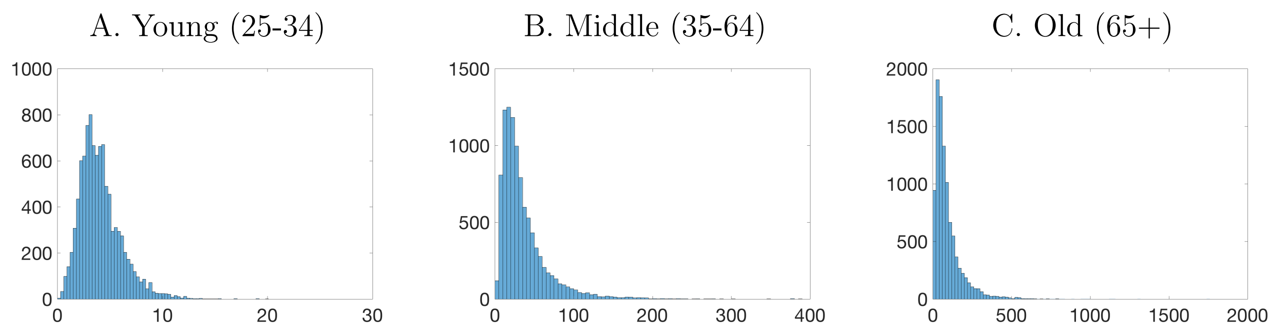

Notes: This figure shows the histograms of the asset holding positions for the three age groups. The horizontal axis denotes the asset holding position per individual and the vertical axis denotes the number of individuals.

The old pay for additional consumption by eating into their assets at a higher rate than the middle and young, which they can afford to do since they experience the highest absolute capital gains from lower interest rates. Average labor supply declines for both middle and young households. Young and middle households hold very few assets while the old hold a relatively large number of assets. Most households accelerate their asset holdings about 10 years before they retire.

In Figure 13, for each age group, we divide households into 'wealthy' and 'poor.' Wealthy households are those in the top decile of wealth and the poor are those in the bottom decile. The figure shows consumption of the wealthy old to be more responsive to interest rate shocks than consumption of the poor old. For the young and the middle-aged, in part because they live primarily off of labor income, and because they substitute higher consumption with an increase in leisure, the difference in the consumption response between the top and bottom deciles is fairly small, although there is a bit of a difference between the wealthy and poor young.

Figure 13: Relative Log Consumption Response by Wealthy and Poor

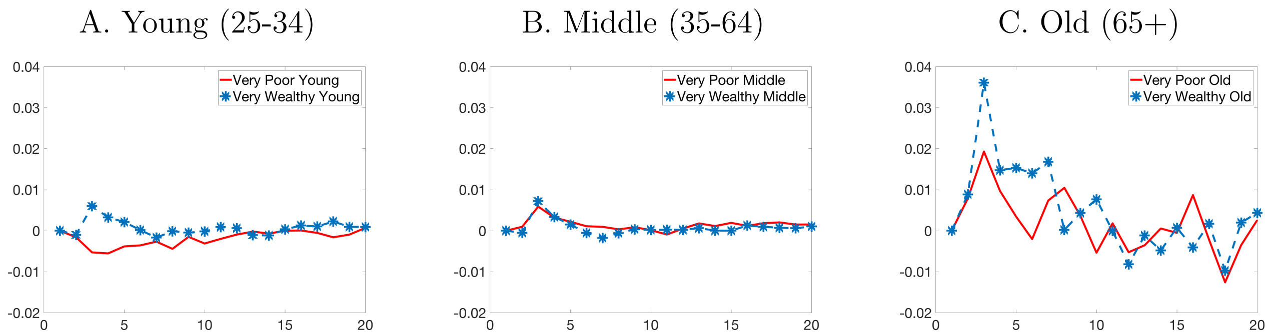

Notes: The figure shows the simulated un-normalized relative responses of mean log consumption across age groups for the wealthy and poor to a -0.16 percent decline in the long-term bond rate, generated by a decline in the short-term rate. If the shock occurs at time $t*$, the relative response for consumption is $\ln(C_t/C_{t*})$. Wealthy households are those in the top decile of asset holdings and poor are those in the bottom decile of asset holdings. The horizontal axis indicates the number of quarters for up to five years after the shock.

6 Conclusion

The empirical evidence shows that old households react more to monetary policy shocks than do middle and young households. This main finding summarizes the weight of evidence gathered across alternative identified monetary policy shocks, empirical methods, and consumption measures. We conjectured four potential features of life-cycle heterogeneity, that together, form the underlying mechanism driving the observed consumption response patterns. They are life-cycle heterogeneity in wealth, portfolio composition, discounting and planning horizons, and labor supply.

We investigate the explanatory power of these ideas with a life-cycle model where households, who face uncertain labor income, death, and interest rates make consumption, saving, and labor supply decisions. The model is able to replicate the most salient features of the data–that consumption of old households are more responsive to monetary policy shocks than younger households.

Understanding potential heterogeneous responses to monetary policy is an interesting topic in its own right. Additionally, as the U.S. population continues to age, our results suggest a potential change in the effectiveness of monetary policy.

References

Anderson, E., A. Inoue, and B. Rossi (2016): "Heterogeneous Consumers and Fiscal Policy Shocks," *Journal of Money, Credit and Banking*, 48, 1877–1888.

Auclert, A. (2016): "Monetary Policy and the Redistribution Channel," *Working Paper, Stanford University.*

Barakchian, S. M. and C. Crowe (2013): "Monetary Policy Matters: Evidence from New Shocks Data," *Journal of Monetary Economics*, 60, 950–966.

Bielecki, M., M. Brzoza-Brzezina, and M. Kolasa (2018): "Demographics, Monetary Policy, and the Zero Lower Bound," *Narodowy Bank Polsk Working Paper No. 284.*

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.

Bunn, P., A. Pugh, and C. Yeates (2018): "The Distributional Impact of Monetary Policy Easing in the UK Between 2008 and 2014," *Staff Working Paper No. 720.*

Carroll, C. D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 23, 61–156.

——— (1997): "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55.

Choi, H., S. Lugauer, and N. C. Mark (2017): "Precautionary Saving of Chinese and U.S. Households," *Journal of Money, Credit and Banking*, 49, 635–661.

Coibion, O., Y. Gorodnichenko, L. Keung, and J. Silvia (2017): "Innocent Bystanders? Monetary Policy and Inequality in the U.S." *Working Paper, UT Austin.*

Di Maggio, M., A. Kermani, and R. Ramcharan (2014): "Monetary Policy Pass-Through: Household Consumption and Voluntary Deleveraging," *Working Paper, Columbia Business School.*

Epstein, L. and S. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Fujiwara, I. and Y. Teranishi (2007): "A Dynamic New Keynesian Life-Cycle Model: Societal Aging, Demographics, and Monetary Policy," *Journal of Economic Dynamics & Control*, 32, 2398–2427.

Gertler, M. and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2017): "Intergenerational Redistribution in the Great Recession," *Working Paper*.

Gornemann, N., K. Kuester, and N. Makoto (2012): "Monetary Policy with Heterogeneous Agents," *FRB Philadelphia WP No. 12-21*.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1, 55–93.

Heathcote, J., K. Storesletten, and G. L. Violante (2010): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118, 681–722.

Jordà, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Krueger, D. and F. Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *The Review of Economic Studies*, 73, 163–193.

Luetticke, R. (2016): "Transmission of Monetary Policy with Heterogeneity in Household Portfolios," *Working Paper, Bonn University*.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106, 3133–3158.

Newey, W. K. and K. D. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703.

Ramey, V. A. (2011): "Identifying Government Spending Shocks: It's All in the Timing," *Quarterly Journal of Economics*, 126, 1–50.

Ríos-Rull, J.-V. and M. Kuhn (2016): "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 1–75.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 1055–1084.

Storesletten, K., C. I. Telmer, and A. Yaron (2007): "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," *Review of Economic Dynamics*, 10, 519–548.

Swanson, E. T. (2018): "Risk Aversion, Risk Premia, and the Labor Margin with Generalized Recursive Preferences," *Review of Economic Dynamics*, 28, 290–321.

Tauchen, G. and R. Hussey (1991): "Quadrature-based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 371–396.

Weil, P. (1989): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421.

Wieland, J. F. and M.-J. Yang (2016): "Financial Dampening," *National Bureau of Economic Research Working Paper No. 22141.*

Wong, A. (2016): "Transmission of Monetary Policy to Consumption and Population Aging," *Working Paper.*

Zeldes, S. P. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97, 305–346.

Appendix (For Online Publication)

A Description of data

Our consumption data is from the quarterly Consumer Expenditure Survey Interview Sample from 1984Q1-2007Q4 compiled by the U.S. Bureau of Labor Statistics. The data are found in the Family Characteristics and Income (FAMILY) files. Each observation is a household in a given time-period.

Our consumption measure is the quarterly expenditure on each component. To measure real consumption expenditures, we deflate each consumption category from the category-specific BLS following Krueger and Perri (2006). Table A–1 reports each consumption category and the BLS CPI code used to deflate each component. Our measure of total consumption is the sum of components 1-19.

Table A–1: CEX consumption categories and CPI categories

	CEX		CPI	
	Category Name	CEX Code (CQ)	Category Name	CPI Code
1	Food	FOODPQ	Food	SAF1
2	Alcohol beverages	ALCBEVCQ	Alcoholic beverages	SAF116
3	Tobacco	TOBACCCQ	Tobacco and smoking products	SEGA
4	Utilities	UTILCQ	Fuels and utilities	SAH2
5	Personal care	PERSCACQ	Personal care	SAG1
6	Household operations	HOUSOPCQ	Household furnishings and operations	SAH3
7	Public transportation	PUBTRACQ	Public transportation	SETG
8	Gas and motor oil	GASMOCQ	Motor fuels	SETB
9	Apparel	APPARCQ	Apparel	SAA
10	Education	EDUCACQ	Tuition expenditures	SEEB
11	Reading	READCQ	Recreational reading material	SERG
12	Health Care	HEALTHCQ	Medical care	SAM
13	Miscellaneous expenditures	MISCCQ	Miscellaneous personal services	SEGD
14	Entertainment	ENTERTCQ	Entertainment	SA6/SAR†
15	House equipment	HOUSEQCQ	Household furnishings and operations	SAH3
16	Vehicles	TRANSCQ-GASMOCQ -PUBTRACQ	Private Transportation	SAT1‡
17	Other lodging	OTHLODCQ	Shelter	SAH1
18	Owned dwelling	OWNDWECQ	Shelter	SAH1
19	Rented dwelling	RENDWECQ	Rent of primary residence	SEHA

Notes: The CPI codes are matched with CEX consumption categories following Krueger and Perri (2006) with the exceptions of entertainment and vehicles.
† Prior to 1998 this is SA6 (Entertainment). From 1998 on it is SAR (recreation).
‡ We combine purchases and vehicle maintenance into Vehicles in the CEX category and use private transport in the CPI.

To calculate the cohort-level measure of consumption used in our analysis, for each grouping on household characteristics (i.e. age and/or income), let $c_{i,t}$ be household consumption divided by the number of household members of household i at time $t =$

$\{1984Q1 - 2007Q4\}$. We follow Anderson et al. (2016) and define consumption by age group as

$$C_t = \left(\frac{1}{H_t} \sum_{i=1}^{H_t} c_{i,t} \right)$$

where H_t is the total number of household observations at a given time t.

B Additional Empirical Results

In this appendix, we report the following additional empirical VAR results.

1. Classification of households into six age groups (25-34, 35-44, 45-54, 55-64, 65-74, 75 and above).
2. Durable goods expenditure responses.
3. Consumption less housing expenditure responses.
4. Alternative consumption averages by age group.
5. Estimating the VAR in levels.
6. Six lags in the VAR.

B.1 Six Age Groups

Here, we consider 6 separate groups: 25-34, 35-44, 45-54, 55-64, 65-74 and 75+. Figures B–1, B–2, and B–3 reports the consumption responses in our VAR from a one standard deviation expansionary HFI-TS, HFI-3, and NG monetary policy shocks, respectively. Once we split the sample up into finer groups, we again see the old respond most to monetary policy shocks. The responses, in general, seem to get progressively stronger as households age.

Figure B–1: Structural VAR – Cumulated Consumption Growth IRF by Age Group to a One Standard Deviation Expansionary HFI-TS Monetary Policy Shock

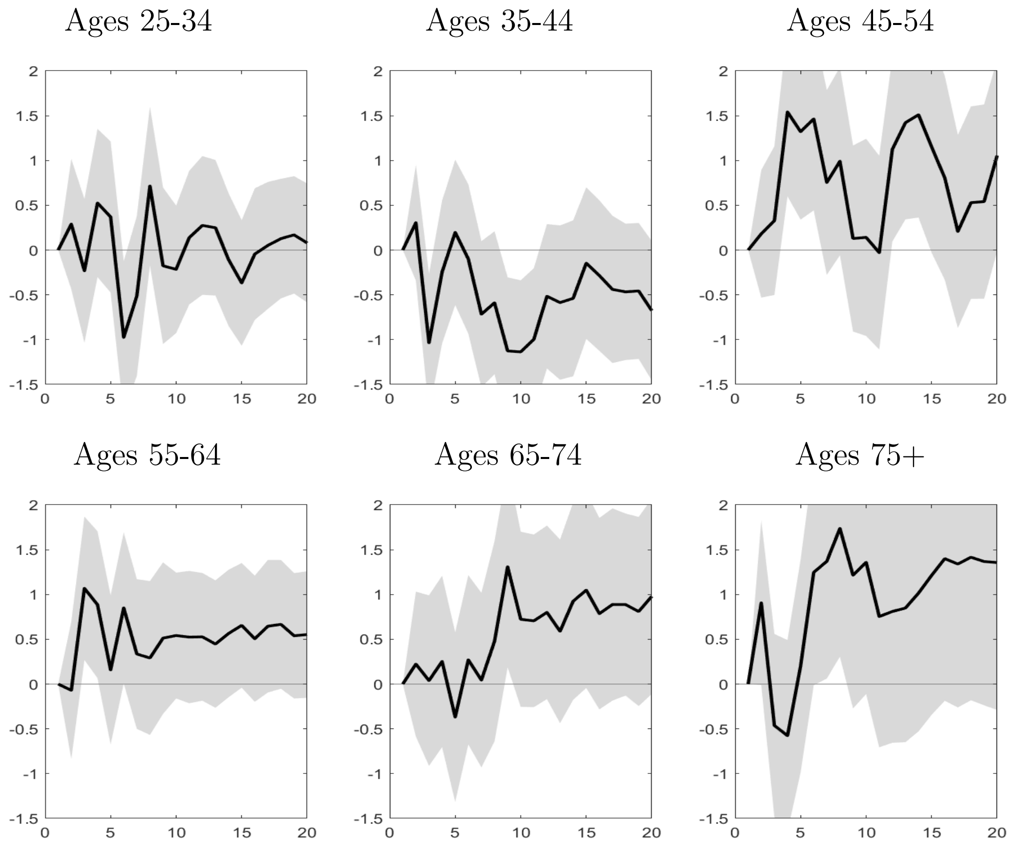

Figure B–2: Structural VAR – Cumulated Consumption Growth IRF by Age Group to a One Standard Deviation Expansionary HFI-3 Monetary Policy Shock

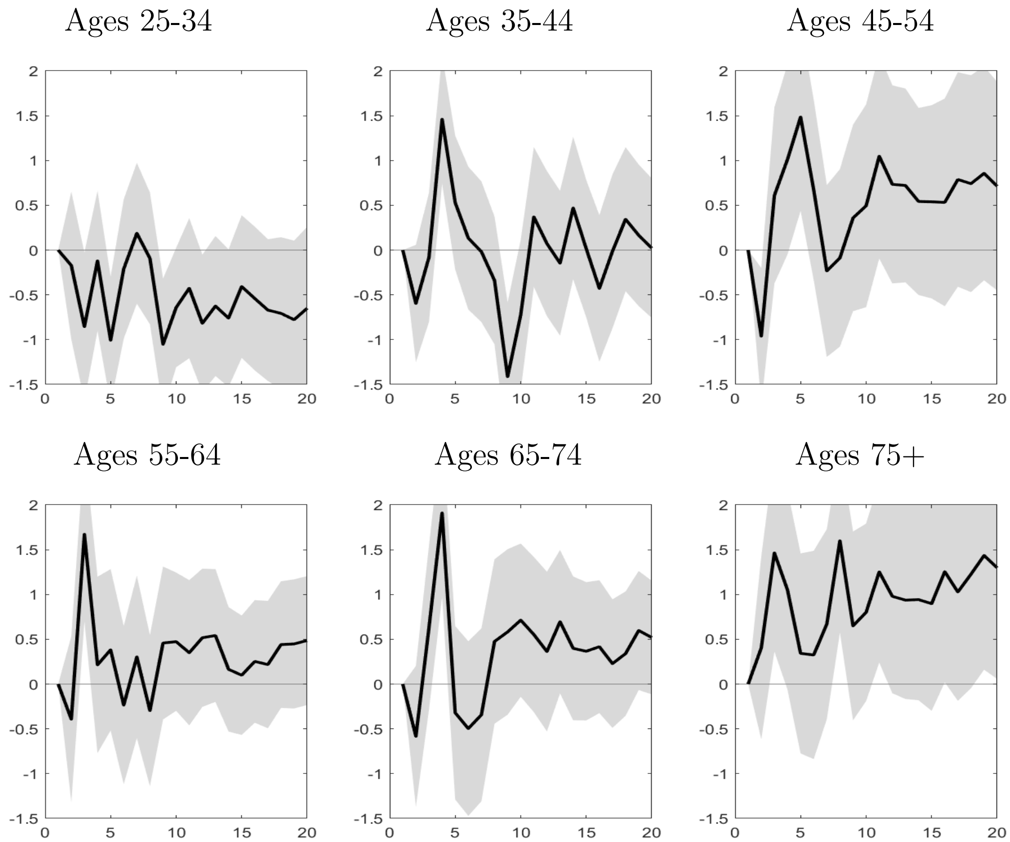

Figure B–3: Structural VAR – Cumulated Consumption Growth IRF by Age Group to a One Standard Deviation Expansionary NG Monetary Policy Shock

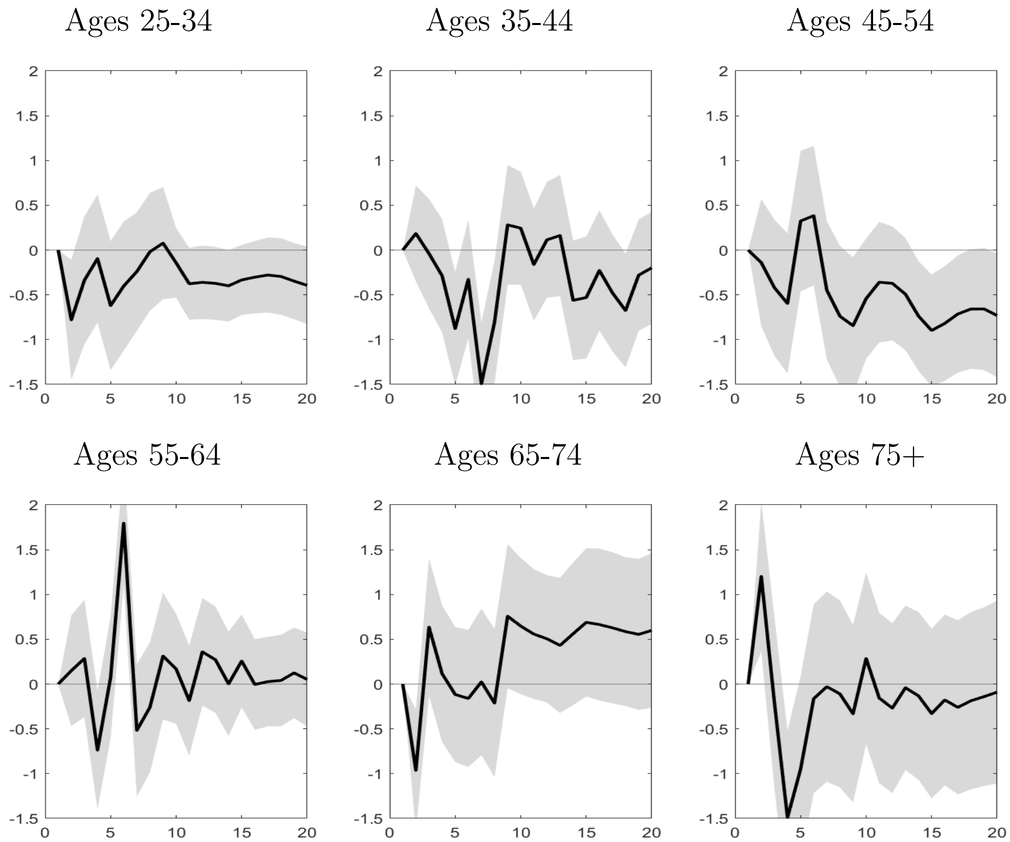

B.2 Durable Consumption

For Figure B–4 we only use responses to consumption on durable goods in our VAR. In the data, the oldest age groups spend a larger share on non-durable consumption than the consumption shown here: at the latter portions of the life-cycle, households already have durable goods acquired through life and are possibly downsizing their ownership of these goods. As a result, the data shows a lot of variation when we sub-divide consumption expenditures to include this smaller component of total consumption.

Figure B–4: Structural VAR – Cumulated Durable Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

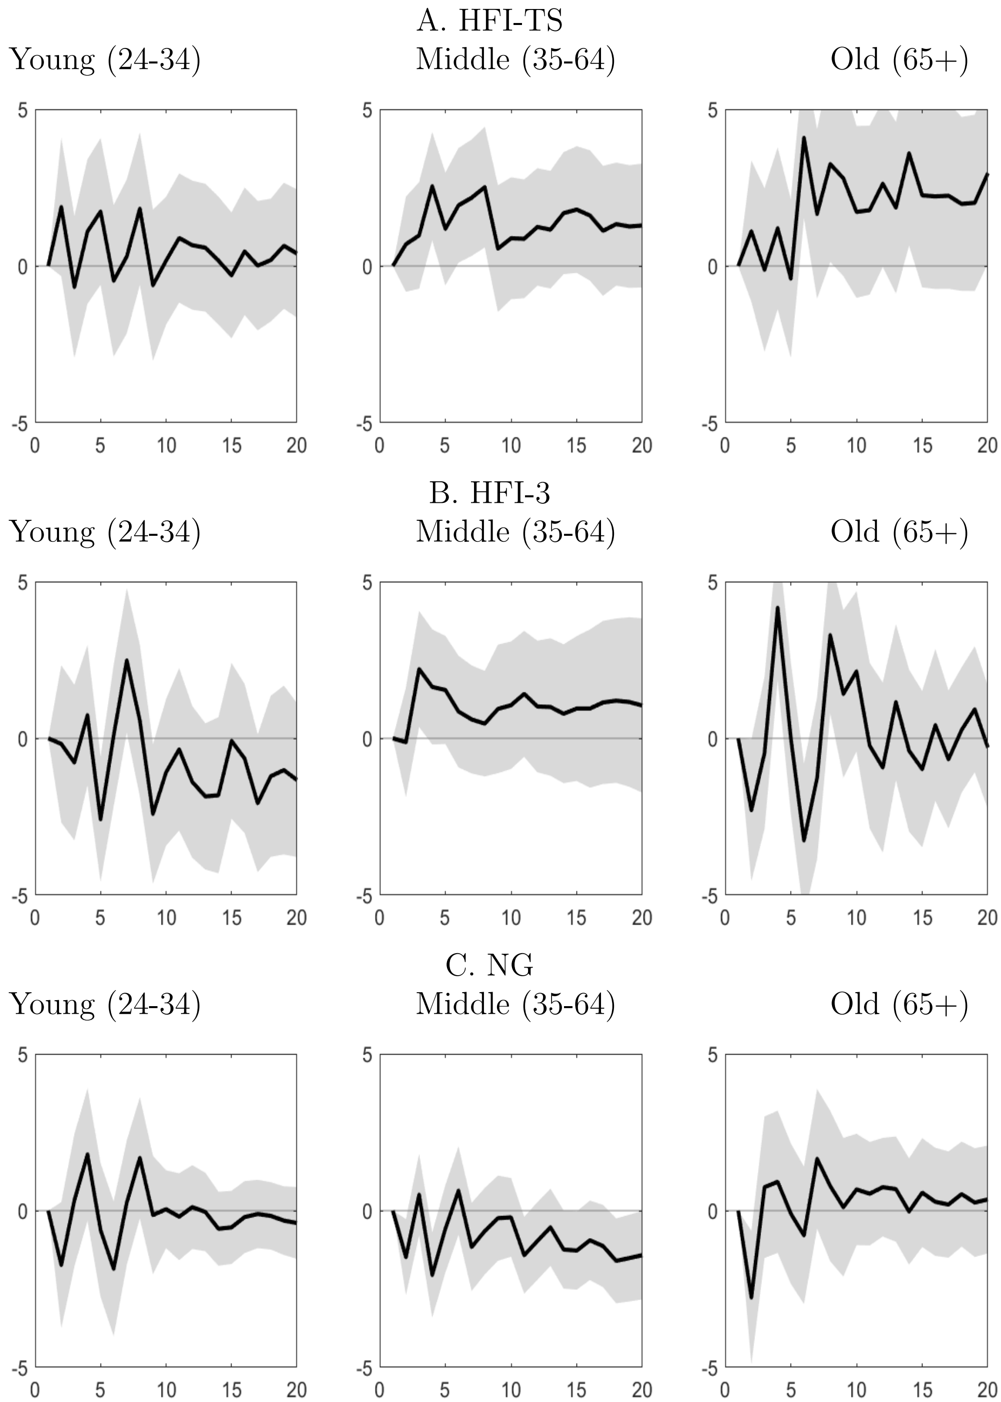

B.3 Consumption Less Housing

Another aspect to consider is the inclusion of housing rents and imputed owner occupied dwelling expenditures. As these categories tend to be rather large, it may be informative to consider consumption expenditures without these categories. We use this consumption measure in our VAR in Figure B–5. These results are closely in line with our baseline results, suggesting housing expenditures are not driving our results.

Figure B–5: Structural VAR – Cumulated Consumption Less Housing Expenditures Growth IRF by Age Group to Expansionary Monetary Policy Shock

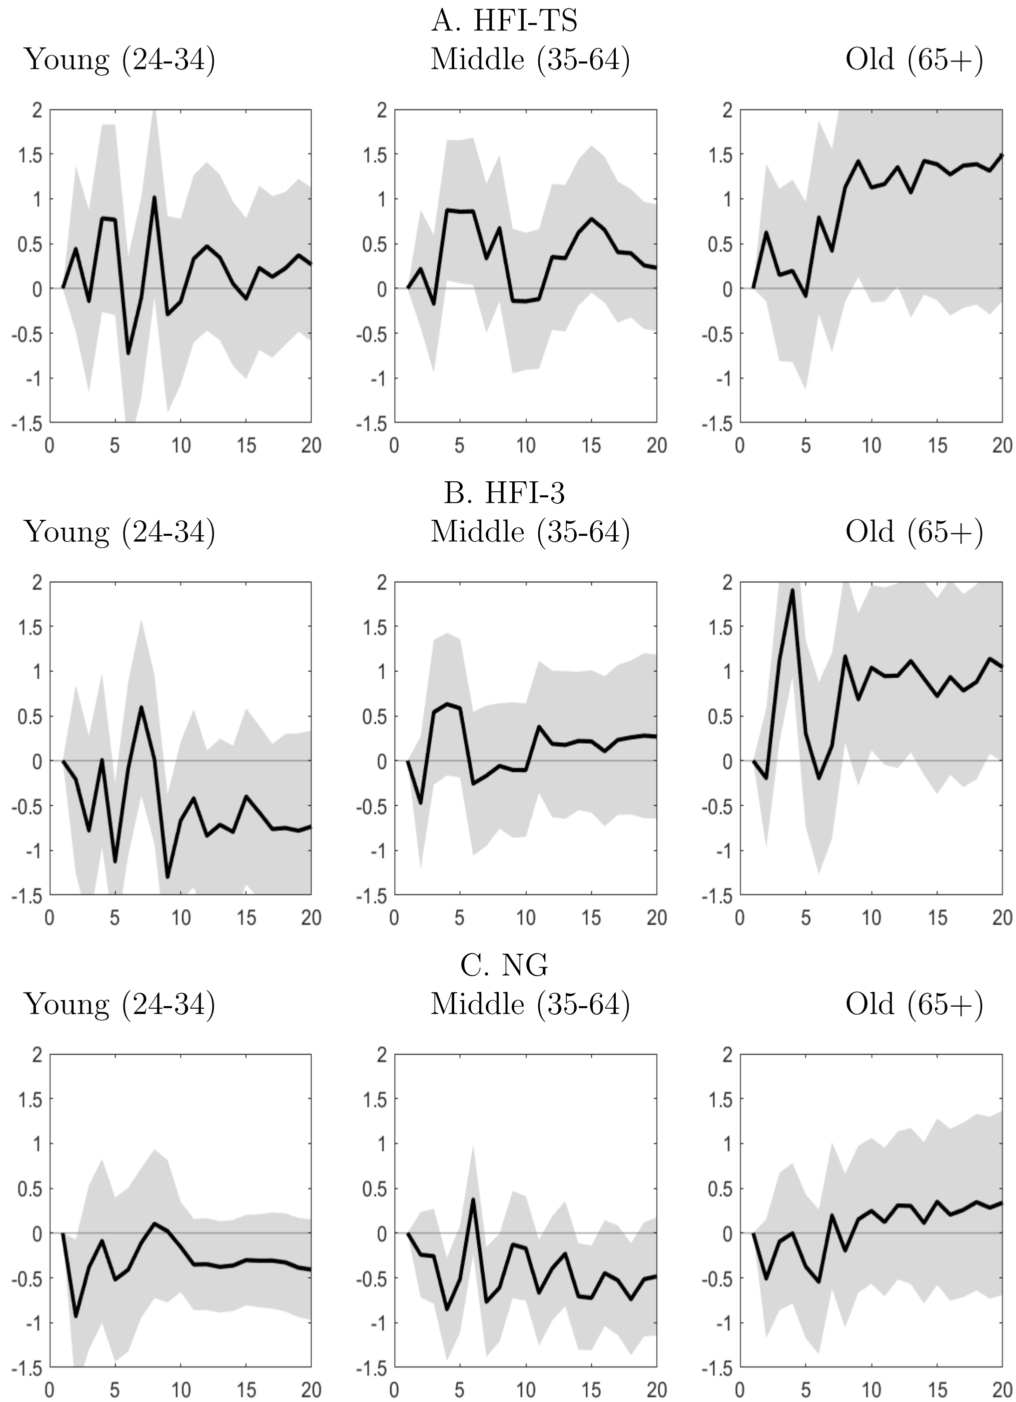

B.4 Alternative Aggregation of Consumption Expenditures by Age

Here, we aggregate consumption by age using the CEX provided sampling weights. Our consumption measure aggregates consumption by age group using the weights, and we transform this into a per-person measure by dividing total consumption for each age group by the weighted cohort sizes. The unattractive aspect of using this measure is that the weights are constructed based on numerous household characteristics to calculate aggregate consumption across all age groups, not just by household age. Nonetheless, Figure B–6 shows the responses to consumption by this measure in our VAR. The results are very close to our baseline specification.

Figure B–6: Structural VAR – Cumulated Consumption (Using CEX Weights) Growth IRF by Age Group to Expansionary Monetary Policy Shock

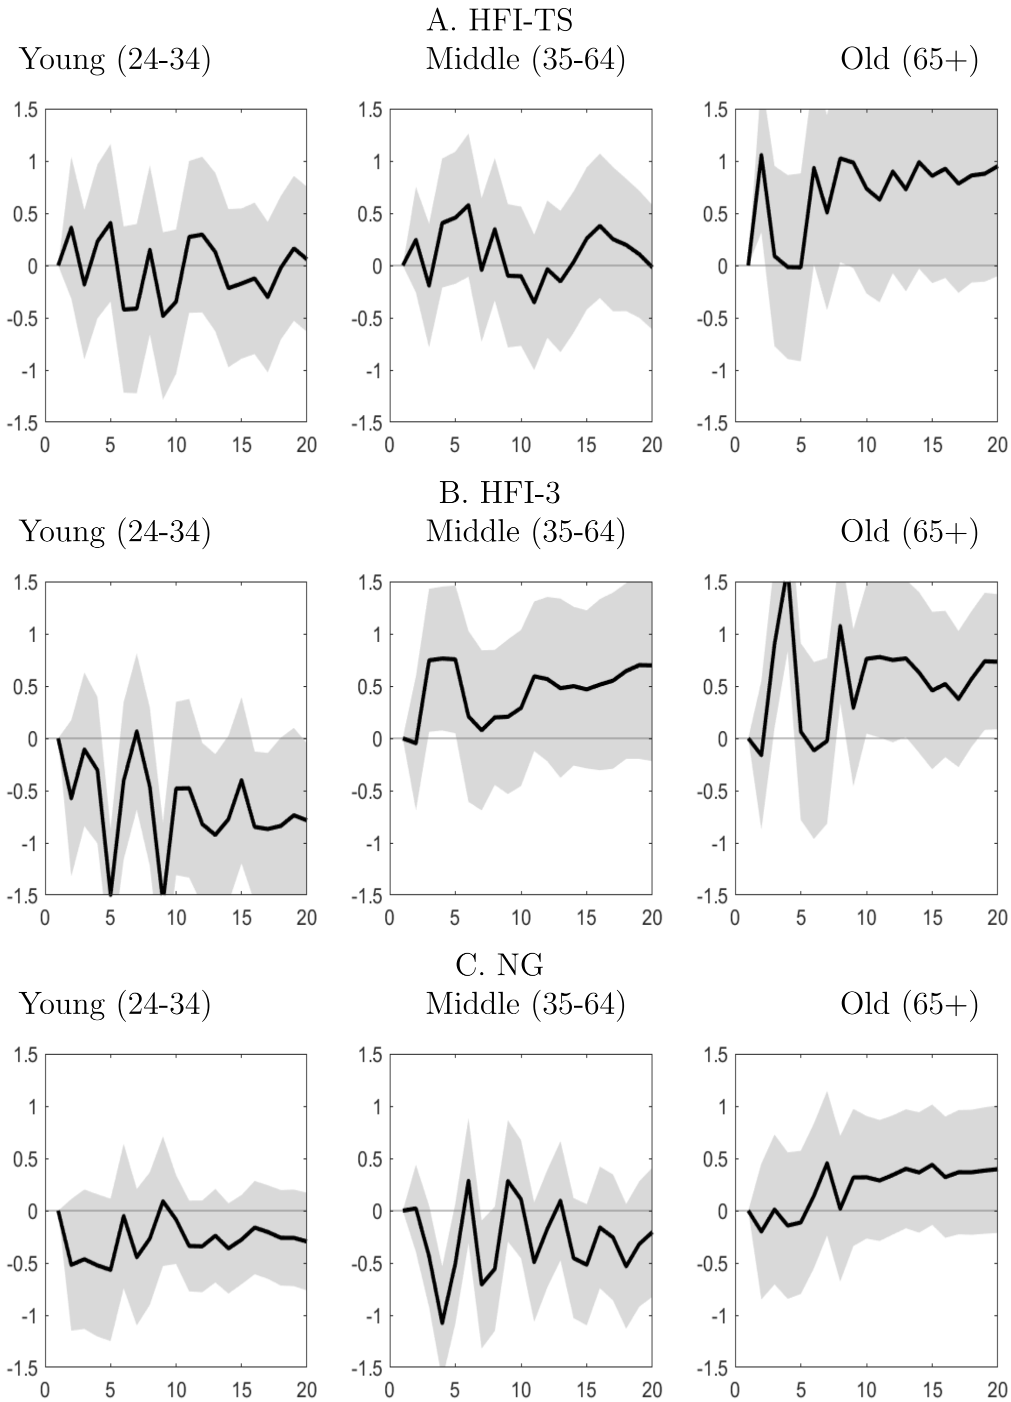

B.5 VAR in Levels

Here, we estimate our VAR in log levels of consumption plotted in Figure B–7. The two main features of this are consistent with our baseline VAR: the old are the most responsive age group and the effects are longest lasting for this age group. Moreover, the level responses for the old are much larger under an innovation in the NG monetary policy shock.

Figure B–7: Structural VAR – Log Level Consumption IRF by Age Group to Expansionary Monetary Policy Shock

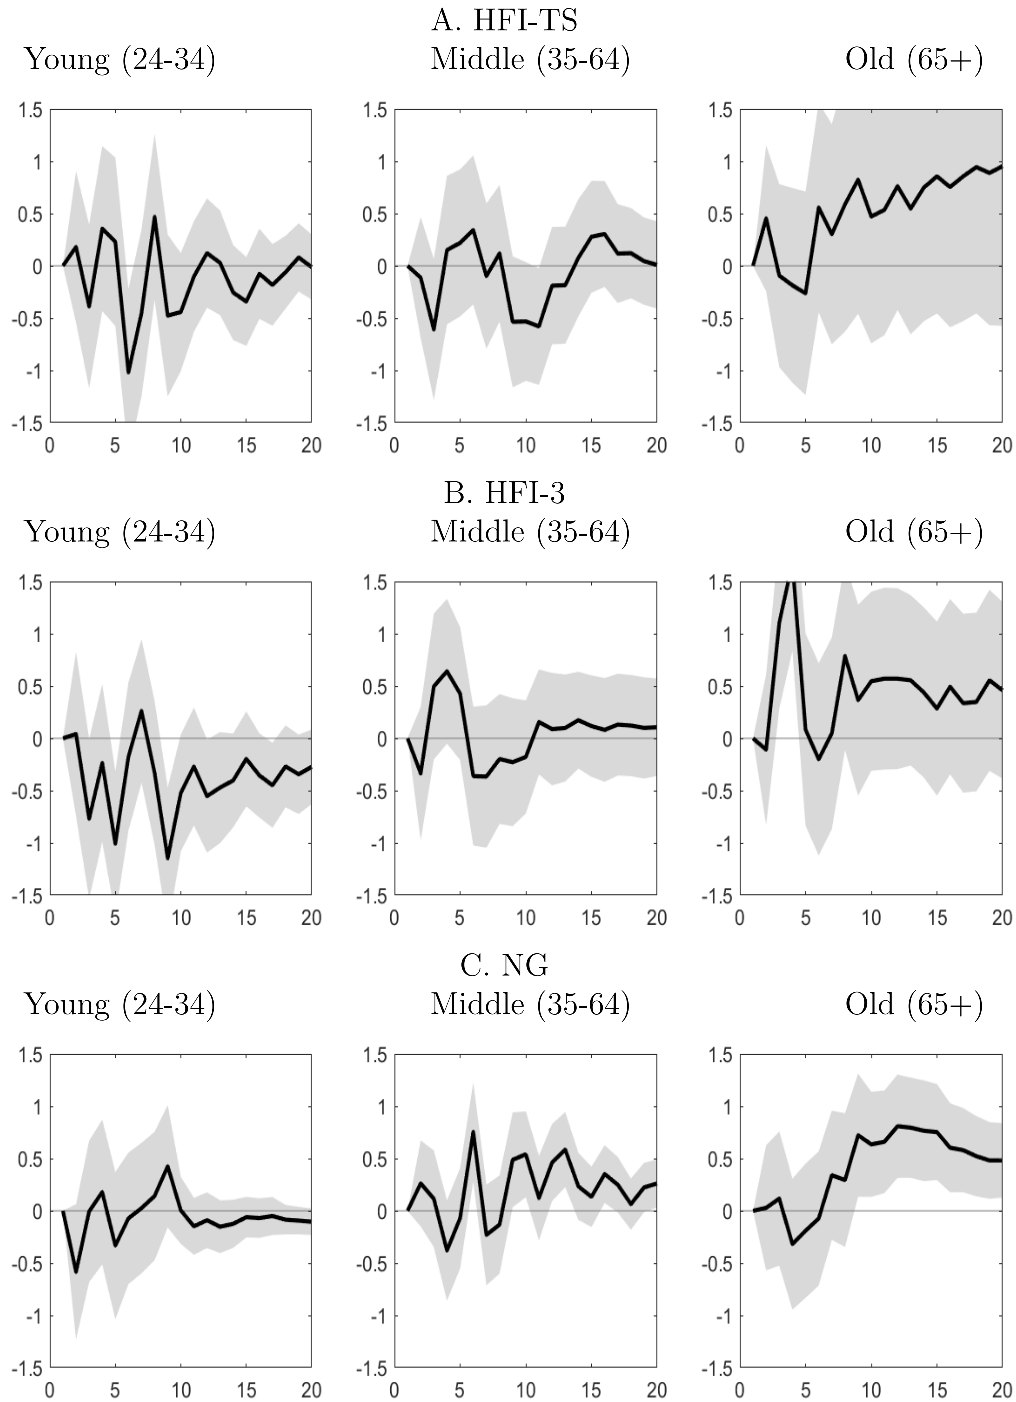

B.6 Two variable VAR

Since the monetary policy shocks are estimated outside of our VAR, it is not dependent of the variables included in our specification. That is, it is appropriate to specify a two-variable VAR with only consumption and the monetary policy shocks. Figure B–8 reports the IRF from an expansionary monetary policy shock using only consumption and the monetary policy shocks. The results are in-line with our main specification, but responses are slightly dampened for the old under the NG monetary policy shock.

Figure B–8: Structural VAR – Cumulated Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock

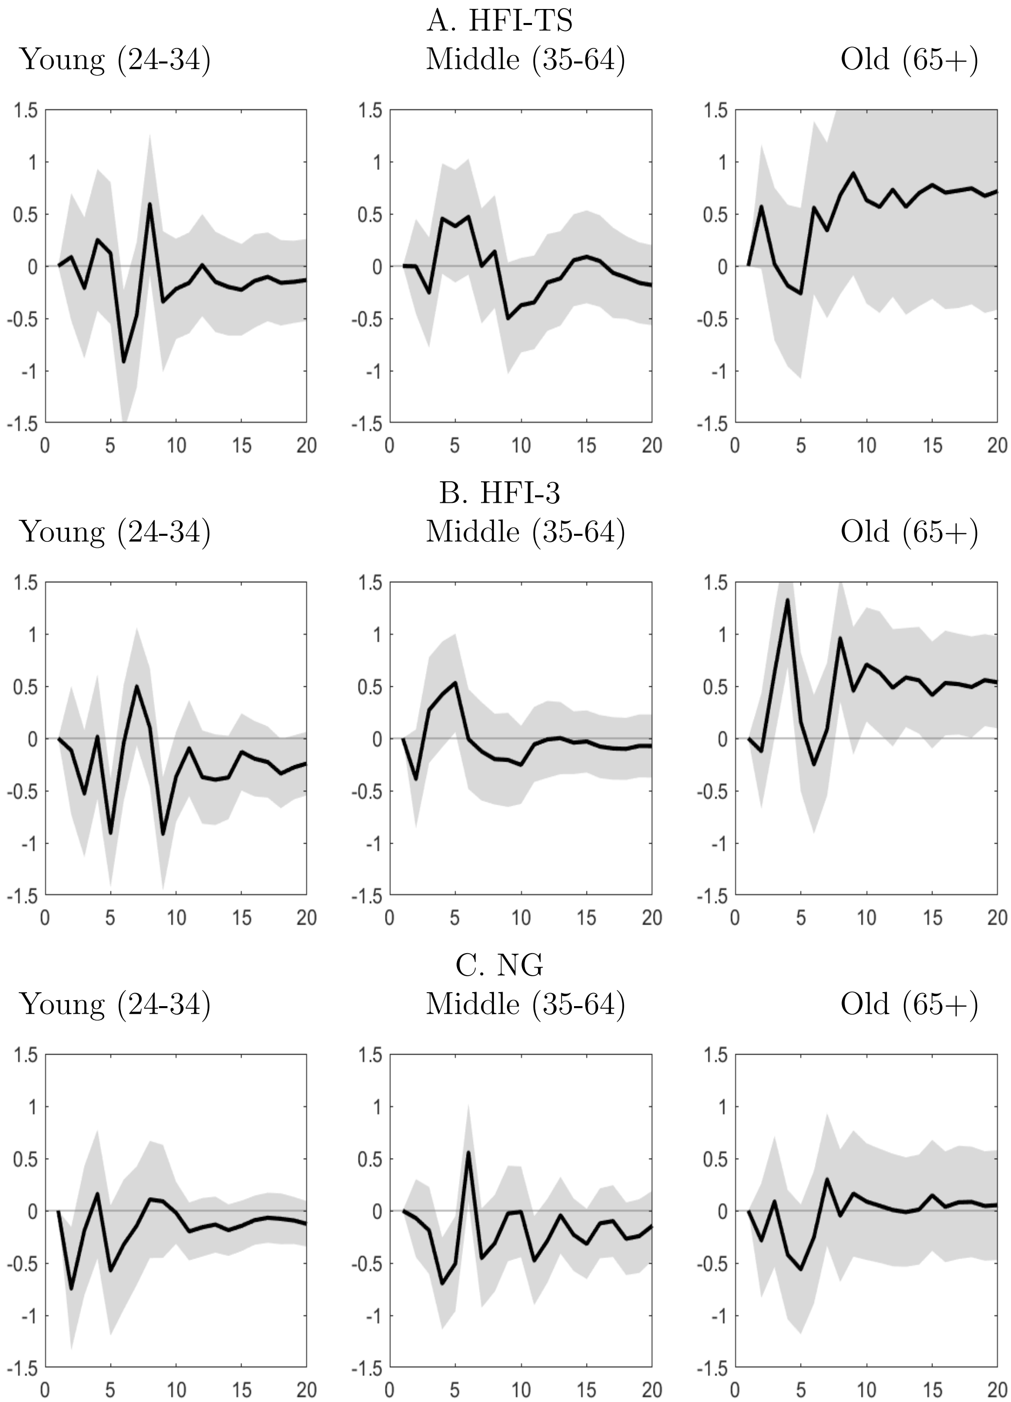

B.7 6 Lags in the VAR

In Figure B–9 we reduce the number of parameter estimates in our VAR by shortening the lags in all variables to 6. The dynamic responses of consumption from the monetary shock with 6 lags produces similar results to our baseline specification.

Figure B–9: Structural VAR – Cumulated Consumption Growth IRF by Age Group to Expansionary Monetary Policy Shock Using 6 Lags

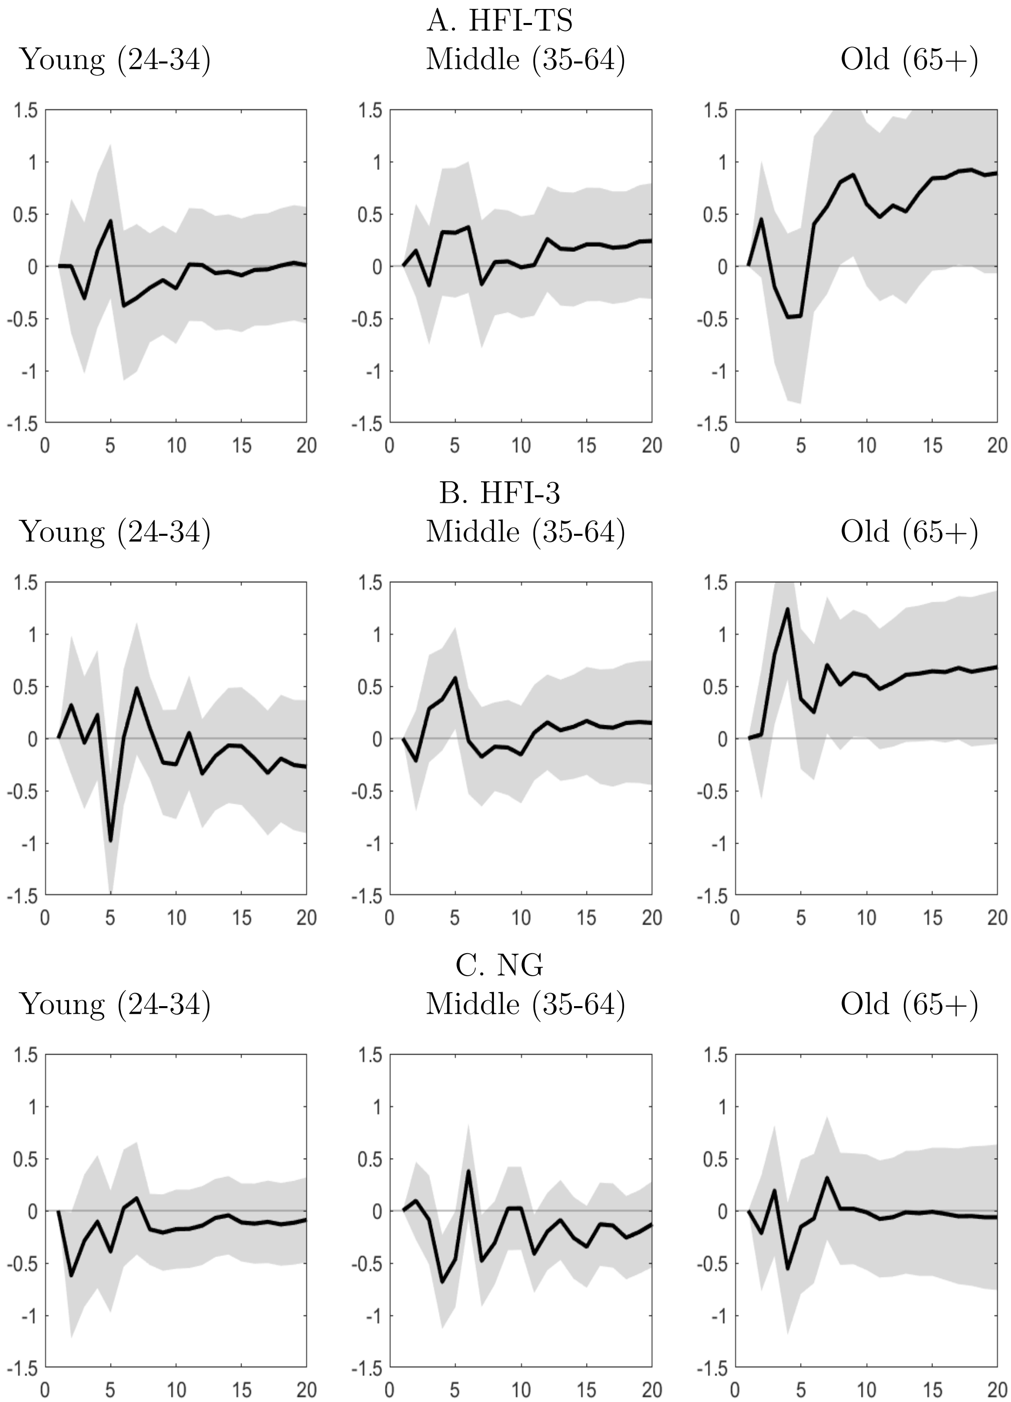

C Stationary Representation of the Model

Because the income process has a unit root, the state space becomes unbounded. To solve the model, we induce stationarity by normalizing income and utility by last period's permanent income. We suppress the individual subscript to avoid clutter. Normalization of the income process follows,

$z < z^*$	$\frac{W_{z,t}}{Y_{z-1,t-1}} = \frac{Y_{z,t}e^{u_t}}{Y_{z-1,t-1}} = M_g G_{z,t} e^{n_{z,t}} e^{u_{z,t}}$	working
$z = z^*$	$\frac{W_{z^*,t}}{Y_{z^*-1,t-1}} = \frac{Y_{z^*,t}}{Y_{z^*-1,t-1}} = p_{rr}$	retirement quarter 1
$z > z^*$	$\frac{W_{z,t}}{Y_{z-1,t-1}} = e^{n_{z,t}} e^{u_{z,t}}$	retirement

Let $\tilde{v}_{z,t} = \frac{V_{z,t}}{Y^{\nu}_{z-1,t-1}}$ and $\tilde{c}_{z,t} = \frac{C_{z,t}}{Y_{z-1,t-1}}$. Normalized utility during the working years is,

$$\tilde{v}_{z,t} = \left\{ (1-\beta) \left(\tilde{c}^{\nu}_{z,t} \left(1 - L_{z,t}\right)^{1-\nu} \right)^{1-\rho} + \beta \left(M_g G_{z,t} e^{n_{z,t}} \right)^{\frac{\nu(1-\rho)}{(1-\gamma)}} \left[E_t \tilde{v}^{1-\gamma}_{t+1} \right]^{\frac{(1-\rho)}{(1-\gamma)}} \right\}^{\frac{1}{1-\rho}}. \tag{16}$$

Let $\tilde{w}_{z,t} = \frac{W_{z,t}}{Y_{z-1,t-1}}$ as defined above and $\tilde{a}_{z,t} = \frac{A_{z,t}}{Y_{z-1,t-1}}$. The normalized budget constraint is,

$$\tilde{a}_{z+1,t+1} M_g G_{z,t} e^{n_{z,t}} = \tilde{a}_{z,t} + r_t \left(\tilde{a}_{z,t} + \tilde{w}_{z,t} L_{z,t} - \tilde{c}_{z,t} \right). \tag{17}$$

In the retirement years, normalized utility and normalized budget constraints are,

$$\tilde{v}_{z,t} = \left\{ (1-\beta) \left(\tilde{c}^{\nu}_{t} \right)^{1-\rho} + \beta \left(M_g G_{z,t} e^{n,t} \right)^{\frac{\nu(1-\rho)}{(1-\gamma)}} \left[E_t \tilde{v}^{1-\gamma}_{t+1} \right]^{\frac{(1-\rho)}{(1-\gamma)}} \right\}^{\frac{1}{1-\rho}}, \tag{18}$$

$$\tilde{a}_{z,t+1} e^{n_{z,t}} = \tilde{a}_{z,t} + r_t \left(\tilde{a}_{z,t} + \tilde{w}_{z,t} - \tilde{c}_{z,t} \right). \tag{19}$$

Adjustments for one-period lived assets. The price of the asset is $P^a_t = (1 + r_t)^{-1}$. The (unnormalized) budget constraints become

$$A_{z,t+1} = (1 + r_t) \left(A_{z,t} + L_{z,t} W_{z,t} - C_{z,t} \right). \tag{20}$$

D Variations on the Model

In this section, we report results from alternative specifications of the model. Figure D–10 shows consumption responses (left panel) from reconfiguring the model such that households can go long or short a one-period bond. The right panel shows the time-path of asset holdings for 2000 individuals over the life cycle. Other model parameters are the same as those in the text, $\gamma = 12, \psi = 0.1, \nu = 0.5$. Size-ordering of the consumption responses goes old > middle $\simeq$ young, but the magnitude of the responses are somewhat smaller than under the long-term asset. With a one-period bond, however, young and middle-aged households are willing to borrow. The right panel shows negative asset holdings in the pre-retirement years. We require non-negative asset holdings after retirement to prevent people from dying in debt.

Figure D–10: Asset is a One-Period Real Bond

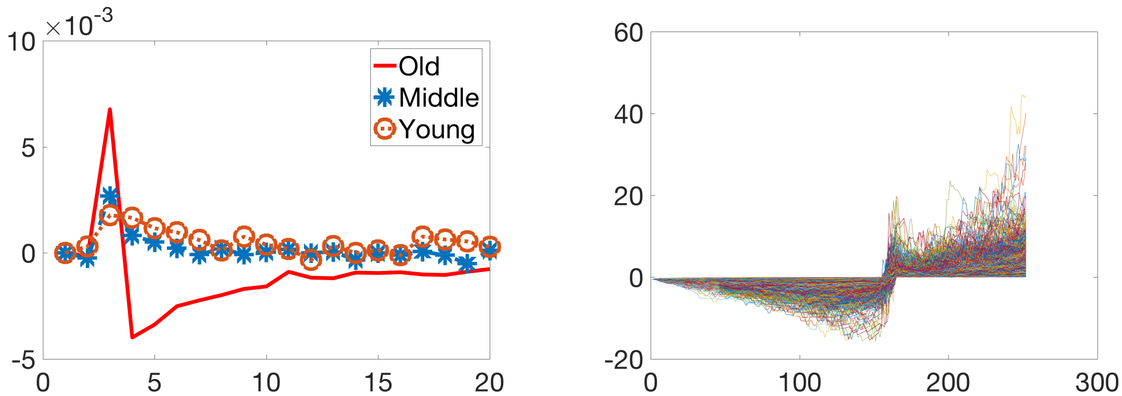

Figure D–11 shows relative consumption responses for the following alternative model specifications

1. Inelastic labor supply

2. Intertemporal elasticity of substitution $\psi = \rho^{-1} = 1.5$.

3. Low risk aversion and low intertemporal elasticity of substitution $\gamma = 2, \psi = 0.1$

4. $\gamma = 4, \psi = 1/4$.

When there is no labor supply decision and income exogenously appears, (Panel A), the size ranking of consumption responses is consistent with our earlier results. Young consumption becomes the most responsive, and middle and old consumption responses are much smaller and similar to each other.

In Panel B, a high intertemporal marginal rate of substitution (on par with values typically assumed in asset pricing studies) induces young consumption to be almost as responsive as old consumption. Both consumptions appear to be a bit excessively responsive, however.

In Panel C, we reduce the risk aversion parameter. Old consumption still exhibits the highest degree of responsiveness, but young consumption becomes too responsive. In the absence of labor supply decisions Panel D would be the case of constant relative risk aversion. Old consumption remains the most responsive, followed by young and middle consumption.

Figure D–11: Alternative Preference Parameter Settings