

NBER WORKING PAPER SERIES

DEMOGRAPHICS AND MONETARY POLICY SHOCKS

Kimberly A. Berg  
Chadwick C. Curtis  
Steven Lugauer  
Nelson C. Mark

Working Paper 25970  
<http://www.nber.org/papers/w25970>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
June 2019

We thank Ben Pugsley and Jim Fackler for comments on an earlier draft and seminar participants at the Robins School of Business, Miami University, University of Kentucky, and the Midwest Macroeconomics Conference. Kimberly Berg received a Miami University Department of Economics Summer 2017 Research Grant to support this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Kimberly A. Berg, Chadwick C. Curtis, Steven Lugauer, and Nelson C. Mark. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Demographics and Monetary Policy Shocks

Kimberly A. Berg, Chadwick C. Curtis, Steven Lugauer, and Nelson C. Mark  
NBER Working Paper No. 25970

June 2019

JEL No. E0,E21,E52,J1,J11

**ABSTRACT**

We decompose the response of aggregate consumption to monetary policy shocks into contributions by households at different stages of the life cycle. This decomposition finds that older households have a higher consumption response than younger households. Amongst older households, the consumption response is also increasing in income. This, along with data on age-related net wealth, presents evidence for a wealth effect playing a role in driving the response patterns. This mechanism is studied further in a partial-equilibrium life-cycle model of consumption, saving, and labor-supply decisions. The model qualitatively explains the 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.

Kimberly A. Berg  
Department of Economics  
Miami University  
Oxford OH  
bergka@miamioh.edu

Chadwick C. Curtis  
Department of Economics  
University of Richmond  
Richmond, VA 23173  
ccurtis2@richmond.edu


Steven Lugauer  
Department of Economics  
University of Kentucky  
225F Gatton  
Lexington, KY 40506  
stevenlugauer@gmail.com

Nelson C. Mark  
Department of Economics  
3060 Jenkins-Nanovic Hall  
University of Notre Dame  
Notre Dame, IN 46556  
and NBER  
nmark@nd.edu

# 1 Introduction

This paper studies how consumption expenditures of different age groups respond to monetary policy shocks. We decompose the response of aggregate consumption to monetary policy shocks into contributions from households at different stages of the life cycle. Empirically, we find that monetary policy shocks have a larger impact on consumption expenditures of older households. Aggregate consumption responses to monetary policy shocks of older households also increases with income. The data show that older households have higher net-wealth than younger households. Since income is also correlated with wealth, this evidence points to a wealth-effect 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 rapid aging of the 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, in demographic time, rapidly increasing. The ratio, which had been fairly steady around 0.2 from 1980 to 2010, is projected to double by 2050 as post-war baby boomers age into retirement.

Age-related heterogeneity in consumption responses could potentially alter the effectiveness of monetary policy as the population ages.

In the paper, Section 2 gives an overview of the data and Section 3 reports the main empirical analyses. We employ impulse responses from structural vector autoregressions (VARs) and from local projections (Jordà (2005)). Using consumption data from the U.S. Consumer Expenditure Survey (CEX), we classify households into young (household head aged 25-34), middle (35-64), and old (65+), and study how aggregate within-group consumption responds to four alternative monetary policy shocks. The U.S. monetary policy shocks were identified and constructed by other researchers. Three of the policy shocks were constructed using the high-frequency identification methods of Barakchian and Crowe (2013), Gorodnichenko and Weber (2016), and Gürkaynak et al. (2005). The fourth uses the narrative/Greenbook methodology of Romer and Romer (2004).<sup>1</sup> Our general finding across two empirical methods and four identified shock series is that old households have the highest proportionate consumption response to monetary policy shocks. Young household consumption appears to be more responsive than middle-aged households, but the evidence here is less definitive.

Section 4 studies the data further in order to understand the economic mechanism that underlies these data patterns. We conjecture four potential life-cycle related effects. 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. Using data from the Survey of Consumer Finances (SCF), we show 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, younger households tend to finance consumption with labor income. They can adjust their labor supply and substitute leisure for consumption in response to monetary policy shocks. Their consumption will be less interest-rate sensitive than older (and retired) households who do not have a labor supply margin and rely on assets whose value are interest-rate sensitive to pay for consumption. 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. Empirically, we explore the wealth mechanism by employing household income as a proxy for wealth. When

---

<sup>1</sup>See Nakamura and Steinsson (2018a) and Nakamura and Steinsson (2018b) for more on estimating monetary policy shocks.

we estimate the structural VARs and local projections for consumption on households classified by income and by age, we observe the highest consumption responses to monetary policy shocks to be by the old high income group.

Our focus is on measuring the aggregate within-age group consumption response because this directly translates into understanding the effects of monetary policy shocks on the aggregate economy. We want to underscore that our aggregative empirical approach is fundamentally different from analyses of the average household response undertaken by research such as Cloyne et al. (2018) and Wong (2018). In Section 5, we contrast these two approaches and discuss how estimated aggregate effects and average household effects can diverge. Wong (2018), for example, obtains a high average response by young households and a low average response by older households. She uses this evidence to motivate the importance of housing in monetary policy transmission. Our own average household response estimates continue to show the old to have the highest responsiveness. We also estimate the average response by the young to be higher than the aggregate young response. If responsiveness is not symmetrically distributed across households, there is no reason for the average and the aggregate estimates to coincide. We show how, if negative policy shocks induce a relatively small number of high-consumption households, within an age group, to increase consumption while the relatively large number of low-consumption households are unresponsive or reduce consumption, the impact on total consumption growth will be larger than the average household effect.

Drawing on our empirical evidence on wealth effects, Section 6 undertakes a complementary analysis for how age and wealth heterogeneity can drive these consumption response patterns using 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 can borrow or lend during their working years by taking short or long positions in a long-term asset but are not allowed to die with negative net worth. 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.

Section 7 concludes the paper. The appendix describes the construction of our micro and

macro data sets, additional robustness analyses, and additional information about the model.

Our paper is part of the growing interest in the macroeconomic implications of agent heterogeneity. Studies of monetary policy transmission with heterogeneity include Gornemann et al. (2012), McKay et al. (2016), and Luetticke (2016). Coibion et al. (2017) studies how monetary policy shocks affect inequality in the United States while Bunn et al. (2018) do so for the United Kingdom. Leahy and Thapar (2019) find that personal income and employment is less responsive to monetary policy shocks in U.S. states with higher percentages of young people. Fujiwara and Teranishi (2007) embed life-cycle behavior in a New Keynesian model, and, similarly, Bielecki et al. (2018) consider how demographic change affects the interest rate within a New-Keynesian model of monetary policy. Also, Doepke and Schneider (2006) find large responses by older households, although their focus is primarily on inflationary episodes as opposed to the identified monetary policy shocks studied in more recent papers.

Research that examines the role of wealth effects for monetary policy include Krueger and Perri (2006), who study the effects on consumption volatility in Italian and U.S. data, Sterk and Tenreyro (2018), who focus on the inflationary consequences of monetary policy for wealth and spending on durables, Glover et al. (2017), who discusses how changes in asset prices during the last recession disproportionately impacted older households, and Auclert (2019), who stresses heterogeneity in the duration of an agents net worth (among other channels). Also, Storesletten et al. (2007) consider idiosyncratic shocks in a life-cycle model, but their focus is on explaining the equity premium puzzle.

Several recent papers research the transmission of monetary policy to consumption through its impact on mortgage finance.<sup>2</sup> Cloyne et al. (2018) study the average effect of monetary policy shocks on non-durable consumption across households with different home ownership categories and find a higher response by renters and those with outstanding mortgages, who tend to be younger than home-owners without mortgages. Di Maggio et al. (2014) study how households of different income levels respond to reductions in mortgage interest payments, while Wong (2018) and Eichenbaum et al. (2018) study how expansionary monetary policy shocks, working through mortgage refinance, changes consumption for younger households.

---

<sup>2</sup>Relatedly, Mian and Sufi (2009), Mian and Sufi (2011), and Mian et al. (2013) stress that changes in housing wealth affect household consumption. Although, Guren et al. (2018) and Guren et al. (2019) argue that housing wealth effects may be moderate in size. The mechanism that we put forth stresses wealth effects across a broad range of assets.

## 2 The Data

This section describes the data used in our main empirical analysis where we estimate the within-age group aggregate consumption response to unanticipated monetary policy shocks. The consumption data is described in Section 2.1 and the monetary policy shocks are discussed in Section 2.2.

### 2.1 Consumption Expenditures

The household consumption expenditure data comes from interview samples of the U.S. Consumer Expenditure Survey (CEX), spanning from 1984Q1 to 2007Q4.<sup>3</sup> We collect quarterly household consumption expenditures on 19 broad categories. We deflate these expenditures by the corresponding categorical price indices from the Consumer Price Index (CPI) to convert into real terms. 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), 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, when long-term asset accumulation generally begins, while people 65 and older are usually retired.

To obtain real per capita household consumption, we divide total real household consumption by the number of household members. The data is not seasonally adjusted so we include seasonal dummy variables in all of the regressions. We do not employ the CEX provided weights because Dynan (2009) (among others) warns that the CEX weights are not justifiable when observations are grouped by demographic characteristics. Appendix A gives a detailed description of the construction of the consumption data. The appendix also reports many additional computations, including showing that our results are robust to the application of the CEX weights.

### 2.2 Monetary Policy Shocks

We consider four alternative measures of identified U.S. monetary policy shocks. These shocks are not our own, but were constructed by other researchers. Using different methodologies and

---

<sup>3</sup>U.S. Department of Labor, Bureau of Labor Statistics, Consumer Expenditure Survey, Interview Survey. Continuous 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.

underlying data, 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 cumulate these monthly observations to a quarterly frequency.

The first measure we consider comes from Barakchian and Crowe (2013) who employ a high frequency identification (HFI) method. Their signal of the policy stance is the term-structure of the federal funds futures contracts for the current month and at 1 through 5 months ahead. The information in the 6 contract horizons is represented by a factor model, and the policy shock is the change in the first factor on the day of an announcement following a Federal Open Market Committee (FOMC) meeting. The change 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 term-structure based shock series as HFI-TRM.

The second monetary policy shock series is the instrument employed in Gertler and Karadi (2015), which they call FF4. They employ the HFI approach of Gürkaynak et al. (2005) and use the change in the three-month ahead federal funds futures price within a 30 minute window of an FOMC announcement. The idea of the short window is to capture that part of the futures price response only to FOMC announcements and not to other news. We refer to this shock series as HFI-3MO.<sup>4</sup>

Our third shock series combines Gürkaynak et al. (2005) from 1990 through 1993, and the monetary policy surprises from Gorodnichenko and Weber (2016) (online appendix Table 16) from 1994 through 2007. Their shock is based on the change in current month federal funds futures within a 60 minute window around FOMC announcements. We label this shock HFI-CMO.<sup>5</sup>

The fourth shock series is constructed following the approach in Romer and Romer (2004). 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 shock is then the residual series, which is argued to be exogenous to current economic conditions and free from anticipatory

---

<sup>4</sup>The HFI-3MO and HFI-TRM shock series are from Valerie Ramey:  
<http://econweb.ucsd.edu/vramey/research.html#data>, accessed in August 2017.

<sup>5</sup>Wong (2018) also uses this combined shock series.



movements. We use the Romer and Romer (2004) updated series by Wieland and Yang (2016) and we refer to the narrative/Greenbook series as NAR-GBK. Their series starts in 1969, but we begin with 1984Q2 to coincide with the beginning of our CEX consumption growth series.

Table 1: Summary Statistics for Alternative Monetary Policy Shocks

A. <u>Mean and Standard Deviation</u>				
	Mean	St. Dev.		
HFI-TRM	-0.001	0.125		
HFI-3MO	-0.045	0.110		
HFI-CMO	-0.048	0.132		
NAR-GBK	0.052	0.277		
B. <u>Correlations</u>				
	HFI-TRM	HFI-3MO	HFI-CMO	NAR-GBK
HFI-TRM	1			
HFI-3MO	0.374	1		
HFI-CMO	0.372	0.778	1	
NAR-GBK	0.308	0.324	0.340	1
C. <u>Autoregressions</u>				
Lag	HFI-TRM	HFI-3MO	HFI-CMO	NAR-GBK
1	-0.076 (-1.017)	0.161 (1.393)	0.071 (0.613)	0.276 (2.920)
2	0.102 (1.075)	0.253 (3.421)	0.084 (0.745)	0.120 (1.075)
3	0.179 (1.934)	0.132 (1.388)	0.246 (2.193)	0.083 (0.847)
4	-0.034 (-0.338)	-0.068 (-0.823)	0.013 (0.114)	0.060 (0.601)
$R^2$	0.043	0.143	0.089	0.154
p-val (Wald)	0.365	0.000	0.138	0.003

Notes: In Panel C, Newey-West t-ratios are in parentheses. The Wald test is for joint significance of the 4-lag coefficients. The starting dates for the series are: 1988Q4 for HFI-TRM, 1990Q1 for HFI-3MO, 1990Q1 for HFI-CMO, and 1984Q2 for NAR-GBK. The ending dates are 2007Q4.

Table 1 reports basic features of the four shock series through 2007Q4. Panel A shows that the shock means are insignificantly different from zero. Also, the standard deviations reveal that the NAR-GBK shocks are about twice the size of the HFI shocks. As one might expect, Panel B shows that HFI-3MO and HFI-CMO are highly correlated with each other. However, the generally low pair-wise correlations with the other shocks points to heterogeneity of information content across the alternative shocks.

Truly exogenous monetary policy shocks should be serially uncorrelated. To check this, Panel C shows fitted fourth-order autoregressions for each of the shocks. The HFI-TRM and

HFI-CMO shocks come closest to satisfying this criteria, as the Wald test for joint significance of lagged coefficients are insignificant for these shocks, while significant for HFI-3MO and NAR-GBK. Although the autoregressions display some evidence against exogeneity for these latter two shock series, we proceed by imposing 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, 0.88 percent per annum, as quarterly changes in the real federal funds rate. 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 main estimation results. In Subsection 3.1, we employ a 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, which includes expenditures on durables and non-durables. Across the four shocks, two estimation methods, and many robustness checks, the weight of the evidence is that the highest consumption response to monetary policy shocks is by older households. Additionally, the heterogeneity across age groups is quantitatively large. The results of several additional robustness checks are reported in Appendix B.

#### 3.1 Structural Vector Autoregressions

The VARs are similar to those employed by Anderson et al. (2016) and Ramey (2011), who study consumption responses to fiscal policy shocks. The three variables in the VAR are, (i)  $g_{c,j,t} = 100\Delta \ln(c_{j,t})$ , the quarterly percentage growth rate of average real per capita consumption of age group  $j = \{\text{young, middle, old}\}$  at time  $t$ , (ii)  $s_t$ , the monetary policy shock and (iii)  $r_t$ , the real federal funds rate.<sup>6</sup> We estimate separate VARs for young (25-34), middle (35-64), and old (65+) households. To avoid clutter in this exposition, we suppress the age group  $j$  subscript.

---

<sup>6</sup>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.

Consumption growth is allowed to respond to contemporaneous monetary policy shocks and changes in the real federal funds rate. The real federal funds rate is affected by contemporaneous policy shocks but not contemporaneous consumption growth. Because the monetary policy shocks are exogenous, neither lags of the shock nor lags of other variables appear in the equation for  $s_t$ . Suppressing the constants, we impose these conditions in the structural VAR as,

$$\begin{pmatrix} 1 & a_{12} & a_{13} \\ 0 & 1 & 0 \\ 0 & a_{32} & 1 \end{pmatrix} \begin{pmatrix} g_{c,t} \\ s_t \\ r_t \end{pmatrix} = \sum_{p=1}^k \begin{pmatrix} b_{p,11} & b_{p,12} & b_{p,13} \\ 0 & 0 & 0 \\ b_{p,31} & b_{p,32} & b_{p,33} \end{pmatrix} \begin{pmatrix} g_{c,t-p} \\ s_{t-p} \\ r_{t-p} \end{pmatrix} + \begin{pmatrix} u_{c,t} \\ u_{s,t} \\ u_{r,t} \end{pmatrix}, \quad (1)$$

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

$$\begin{pmatrix} g_{c,t} \\ s_t \\ r_t \end{pmatrix} = \sum_{p=1}^k \begin{pmatrix} c_{p,11} & c_{p,12} & c_{p,13} \\ 0 & 0 & 0 \\ c_{p,31} & c_{p,32} & c_{p,33} \end{pmatrix} \begin{pmatrix} g_{c,t-p} \\ s_{t-p} \\ r_{t-p} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \epsilon_{3,t} \end{pmatrix}. \quad (2)$$

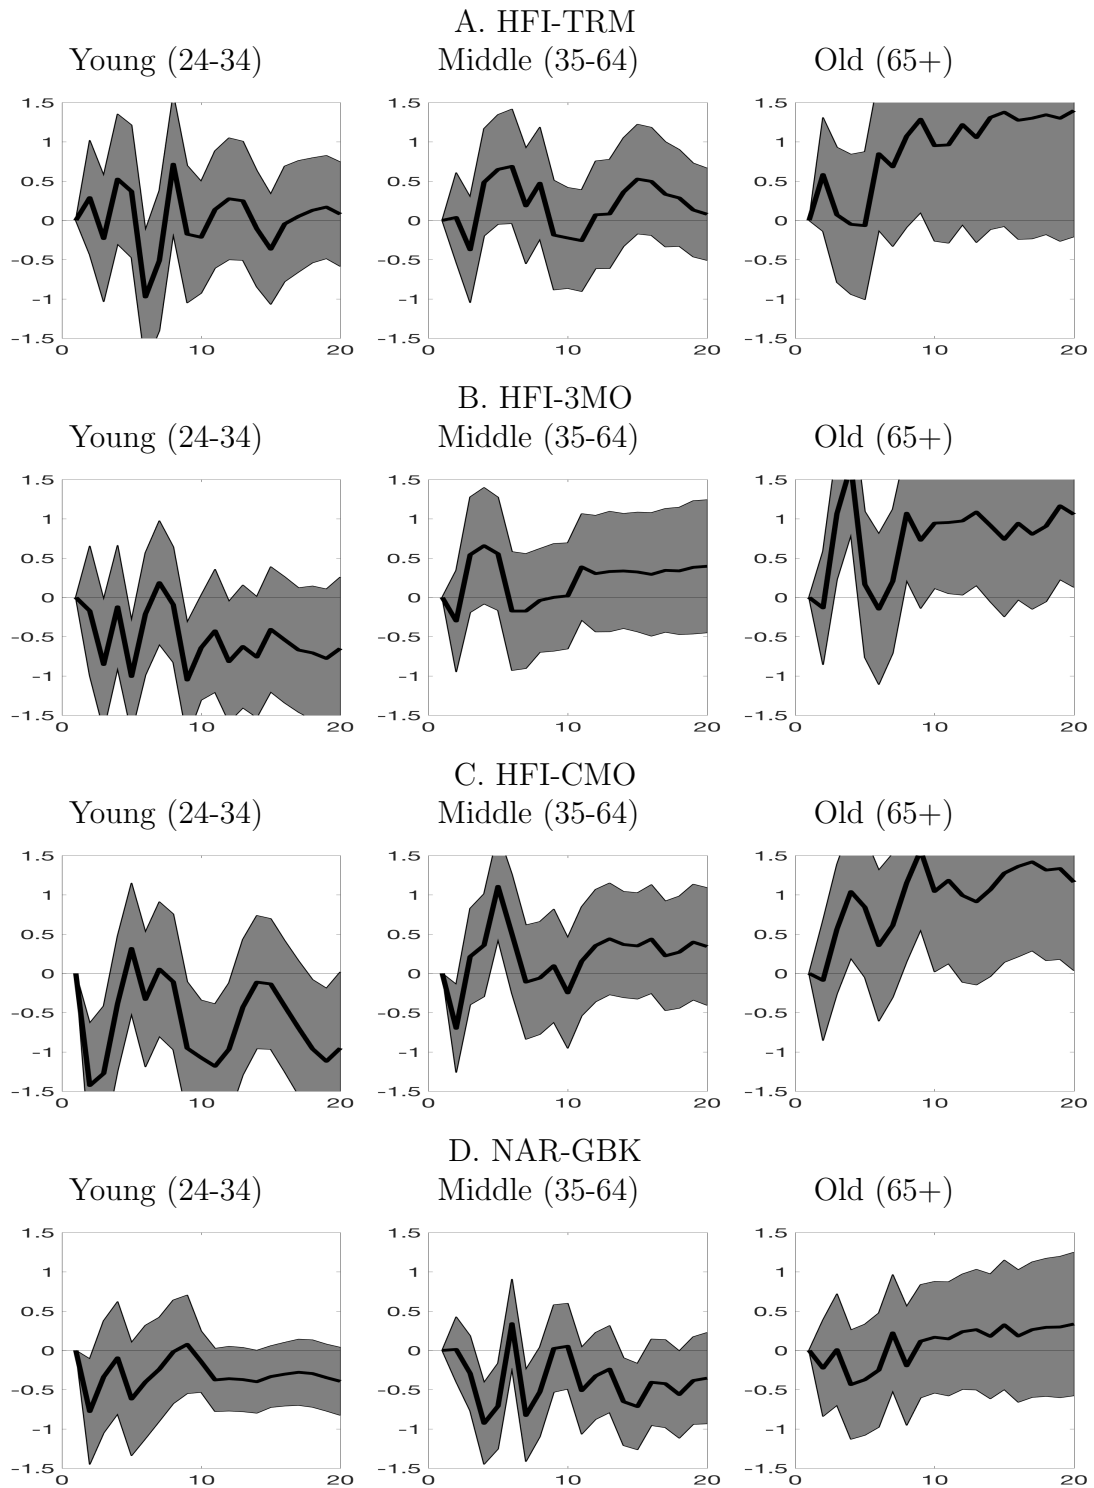
Due to the relatively short time-span of the data, imposing these theoretical restrictions lightens the parameterization of the VAR and preserves degrees of freedom.<sup>7</sup>

We estimate the VARs with  $k = 8$  lags. 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 consumption response in percent. The response to the HFI-TRM shock are shown in Panel A, to the HFI-3MO shock in Panel B, to the HFI-CMO shock in Panel C, and to the NAR-GBK shock in Panel D. 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)).

---

<sup>7</sup>After accounting for 8 start-up values, the HFI-TRM VAR has 69 quarterly observations, HFI-3MO has 64, HFI-CMO has 64, and NAR-GBK has 87.

Figure 2: Structural VAR – Cumulated Consumption Growth Impulse Response 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  $\pm$  one standard error asymptotic confidence bands. The horizontal axis indicates quarters following the shock. The vertical axis is measured in percent.

Following an expansionary HFI-TRM shock, shown in Panel A, the largest consumption response is by the old. Responses by the young and middle age groups are muted in comparison. After about eight quarters, consumption for the old households has increased dramatically, and the effect seems permanent.<sup>8</sup> The peak consumption response for old households is about twice as high as the peak for the middle-aged. As time passes, the impact on young households dissipates.

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

An expansionary HFI-CMO shock, shown in Panel C, also leads to an apparent permanent increase in consumption by old households. The shock induces a relatively large but temporary increase in middle-aged consumption. Consumption of the young declines.

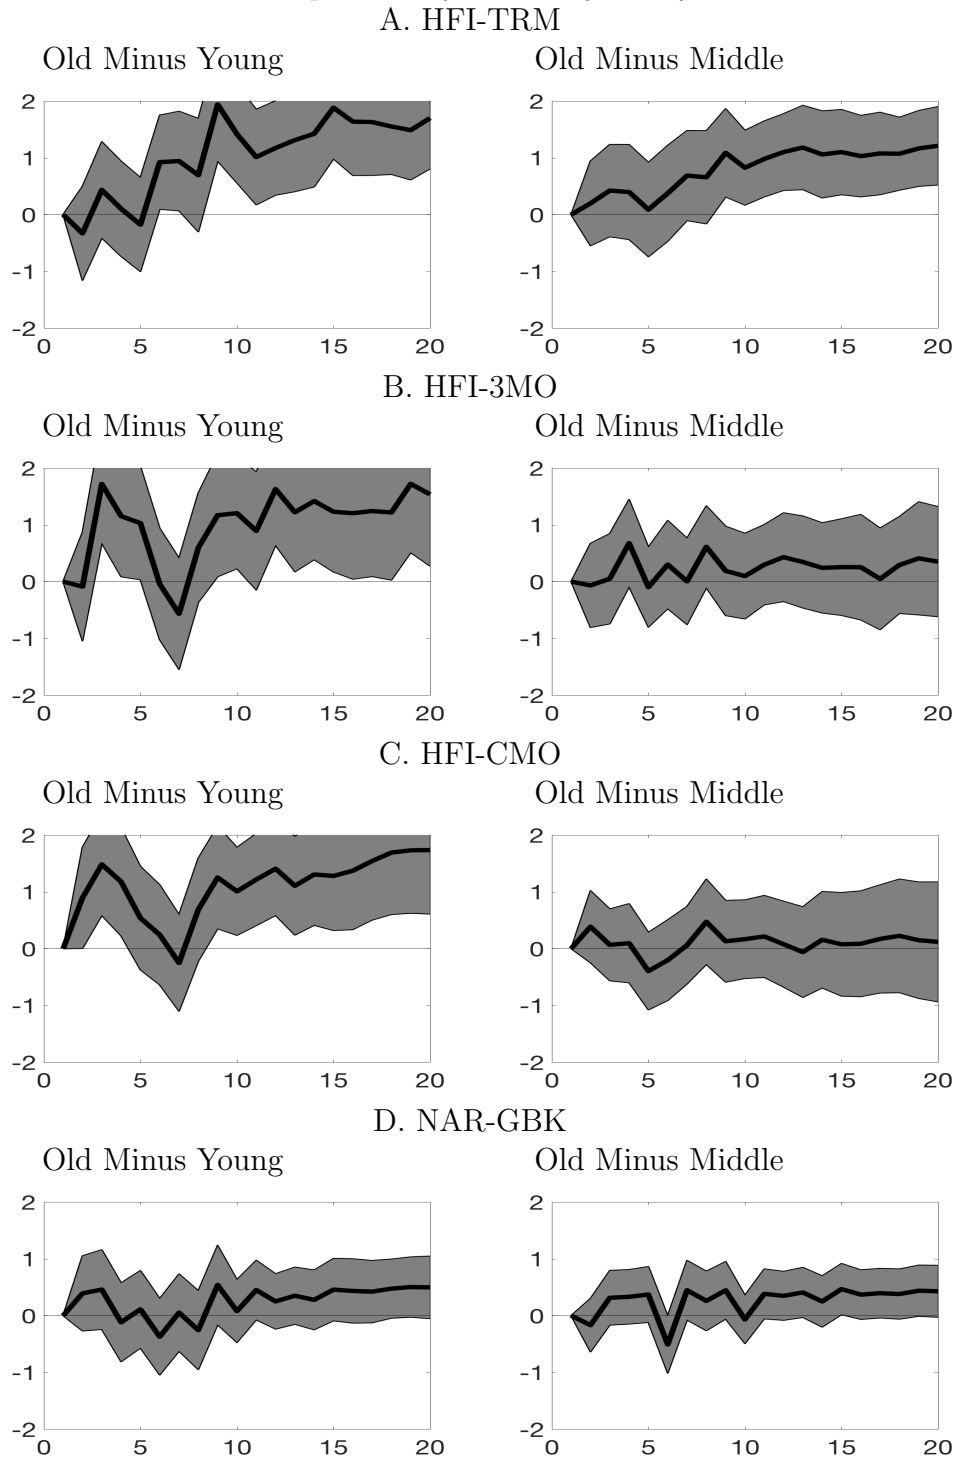
The consumption responses to an expansionary NAR-GBK shock, shown in Panel D, are relatively subdued compared to the responses to the other policy shocks. Young and middle consumption display similarity in timing and magnitudes, both exhibiting modest decreases. Old consumption also initially declines in response to the negative NAR-GBK shock, but then the response gradually turns positive.

To summarize, for each of the four monetary policy shocks, we uncover heterogeneity in consumption responses across age groups with the old having the largest response. Figure 3 plots old minus young and old minus middle impulse responses to draw out the comparison amongst the age groups. The differences across age groups is also quantitatively large. Conditional on the particular policy shock, the old clearly have the most positive response in consumption.

---

<sup>8</sup>Romer and Romer (2004) and Coibion et al. (2017) also find persistence in the effects from monetary policy shocks.

Figure 3: Structural VAR – Cumulated Consumption Growth Impulse Response for Old Minus Young and Old Minus Middle 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  $\pm$  one standard error asymptotic confidence bands. The horizontal axis indicates quarters following the shock. The vertical axis is measured in percent.

*Changing Monetary Policy Effectiveness over Time.* The heterogeneous 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 aggregate our age-specific impulse responses and combine them with alternative demographic profiles to estimate policy effectiveness at different points in time. Let  $c_A$  be real per capita consumption and  $N_A$  be the number of people in the aggregate. Similarly, let  $c_y, c_m, c_o$  ( $N_y, N_m, N_o$ ) be consumption (numbers of) of young, middle, and old respectively. Then aggregate consumption is  $N_A c_A$ , and the approximate relative change in this aggregate is

$$N_A \Delta \ln(c_{A,t}) = N_y \frac{\bar{c}_y}{\bar{c}_A} \Delta \ln(c_{y,t}) + N_m \frac{\bar{c}_m}{\bar{c}_A} \Delta \ln(c_{m,t}) + N_o \frac{\bar{c}_o}{\bar{c}_A} \Delta \ln(c_{o,t}), \quad (3)$$

where estimates of the change in young, middle, and old household real per capita consumption come from the estimated VARs, and the number of young, middle, old, 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. 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. We abstract from changes in family size and composition within-age groups.

We estimate the response of cumulated aggregate consumption growth to each of the four monetary policy shocks using population weights in years 1990, 2010, and 2030. Figure 4 displays the results. 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.<sup>9</sup> These figures highlight how dramatically the age-distribution changes the likely effectiveness of monetary policy, especially for the HFI-TRM, HFI-3MO, and HFI-CMO 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 population aging from 1990 through 2030 increases monetary policy effectiveness when measured by the three ‘high-frequency’ shocks, the differential consumption responses to the NAR-GBK


---

<sup>9</sup>We are calculating the net area between the curves in Panel A and the zero line.


shock are much smaller.

Figure 4: Structural VAR – Monetary Policy Effectiveness

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



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. In Panel A, the horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent in Panels A and B. We estimate the consumption responses by age group to a monetary policy shock and hold them 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.

### 3.2 Local Projections

Using the structural VARs, we find monetary policy shocks to impart heterogeneous consumption responses across households of different age groups. In this section, we examine the robustness



of these results with respect to the empirical procedures by employing local projections (Jordà (2005)) as an alternative strategy. As in the VARs, we control for past consumption growth, past real federal funds rate, and seasonal effects in the regressions. The local projections are the sequence of regressions at horizons  $h = 1, \dots, 20$ , estimated separately for each age group (group subscript suppressed),

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

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. To preserve degrees of freedom, we control for past consumption growth at the yearly horizon instead of including 8 lags of consumption growth.<sup>10</sup>

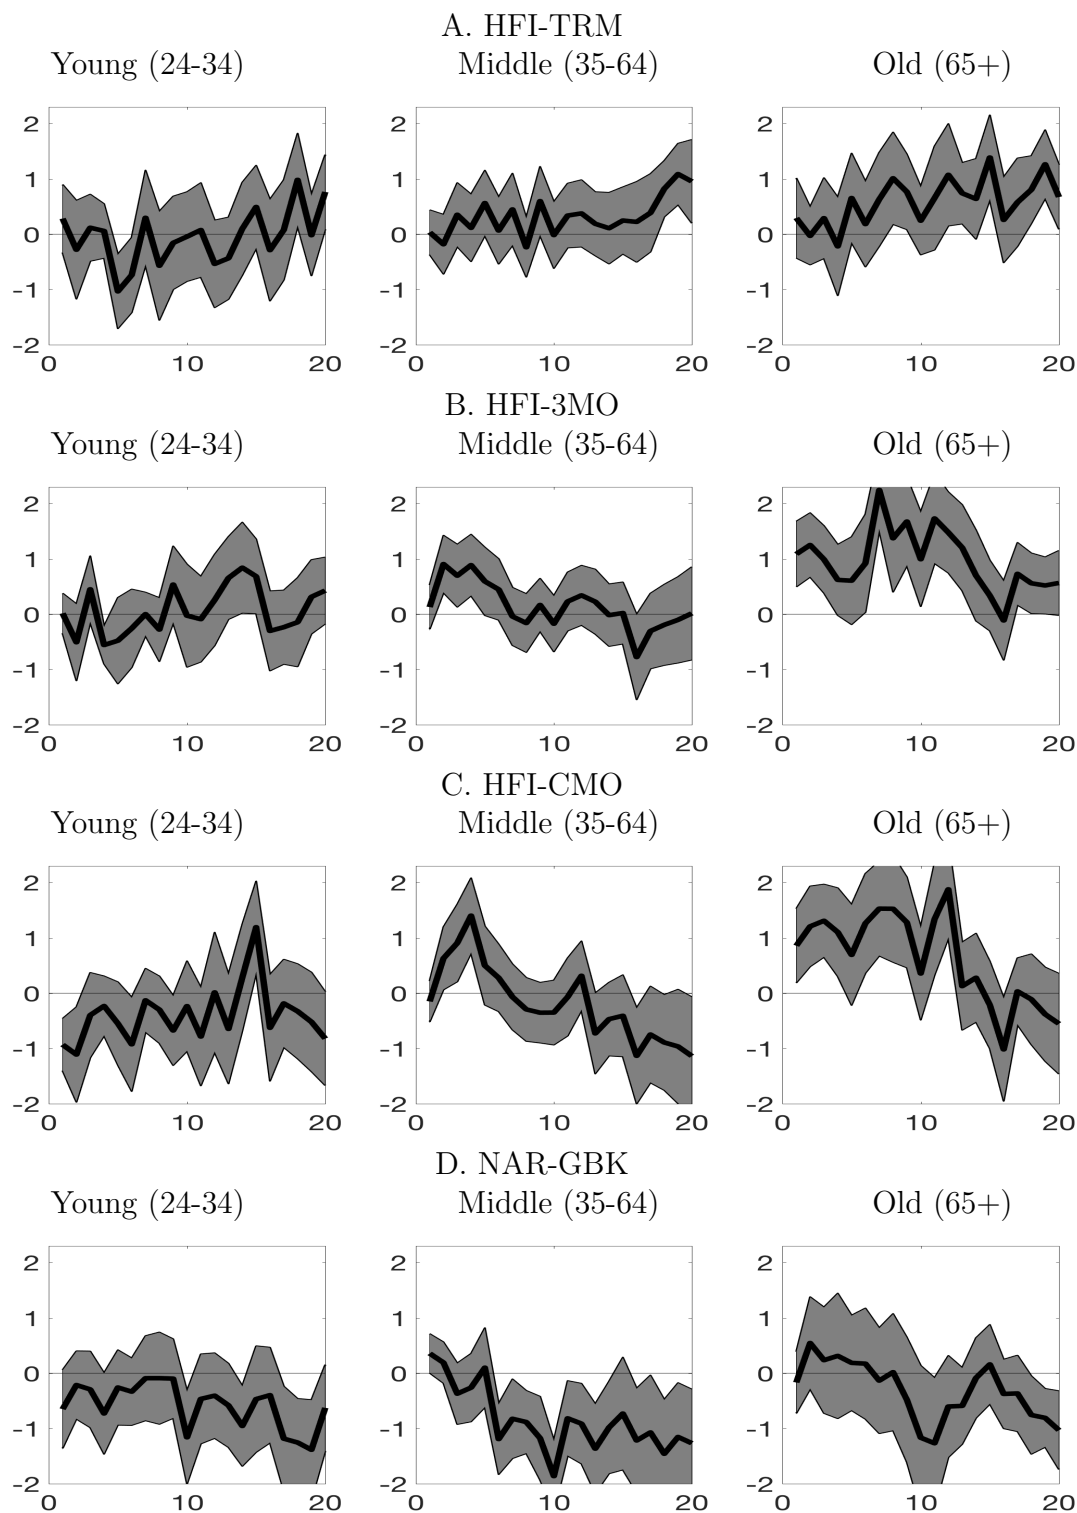
The coefficient of interest is  $\beta_h$ , which measures the percent change in the consumption response from time  $t$  to  $t + h$  due to the monetary policy shock at time  $t$ . To express the response to an expansionary shock, Figure 5 displays plots of  $-\beta_h$  with  $\pm 1$  Newey and West (1987) standard-error bands.

The local projections exhibit consumption response heterogeneity by age group. The response of the old to the HFI-3MO and HFI-CMO shocks are higher than for the younger age groups over much of the following 20 quarters. The old also exhibit a higher consumption response to the other two shocks relative to the middle-aged group. Although, the old consumption increases only modestly for about two years after a NAR-GBK shock, the response then turns negative. The middle group exhibits an initial positive response to the HFI-3MO, HFI-CMO, and NAR-GBK shocks, but it quickly dies out. Looking across the panels, the responses by the young are generally near zero and sometimes quite negative.

---

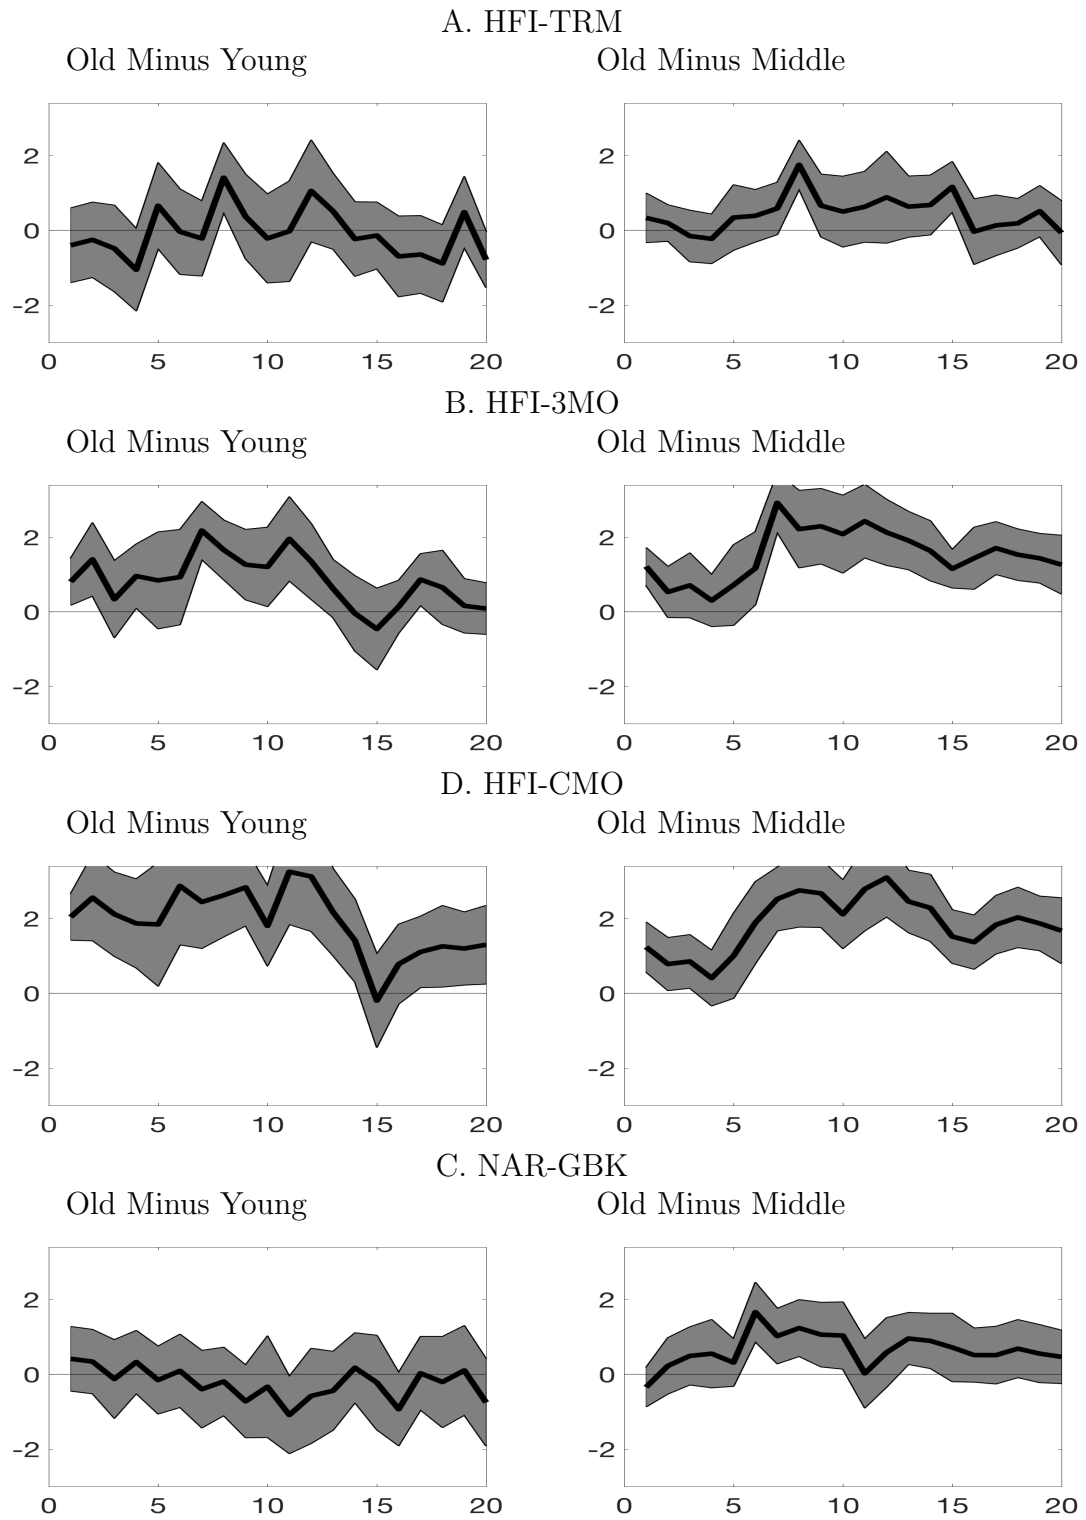
<sup>10</sup>If the shocks  $s_t$  are truly exogenous, as pointed out by a referee, controlling for lagged consumption and the real federal funds rate may not be necessary. In Appendix B, we show that our local projection results are robust to omitting these controls.

Figure 5: Local Projections – Cumulated Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



Notes: This figure plots  $-\beta_h$  from Equation (4). Shaded areas are  $\pm$  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.

Figure 6: Local Projections – Cumulated Consumption Growth Impulse Response for Old Minus Young and Old Minus Middle to Expansionary Monetary Policy Shock



Notes: Shaded areas are  $\pm$  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 local projections results, old consumption responds strongly positively to the HFI-TRM and HFI-3MO expansionary monetary policy shocks and, within the first three years after the shock, responds positively to the HFI-CMO shock. To draw out the differences across age groups, Figure 6 plots the difference between the old and young response and the old and middle response. Old consumption generally increases the most. The relative ranking of consumption responses between young and middle households is less definitive.

### 3.3 Response of Non-Durable Consumption

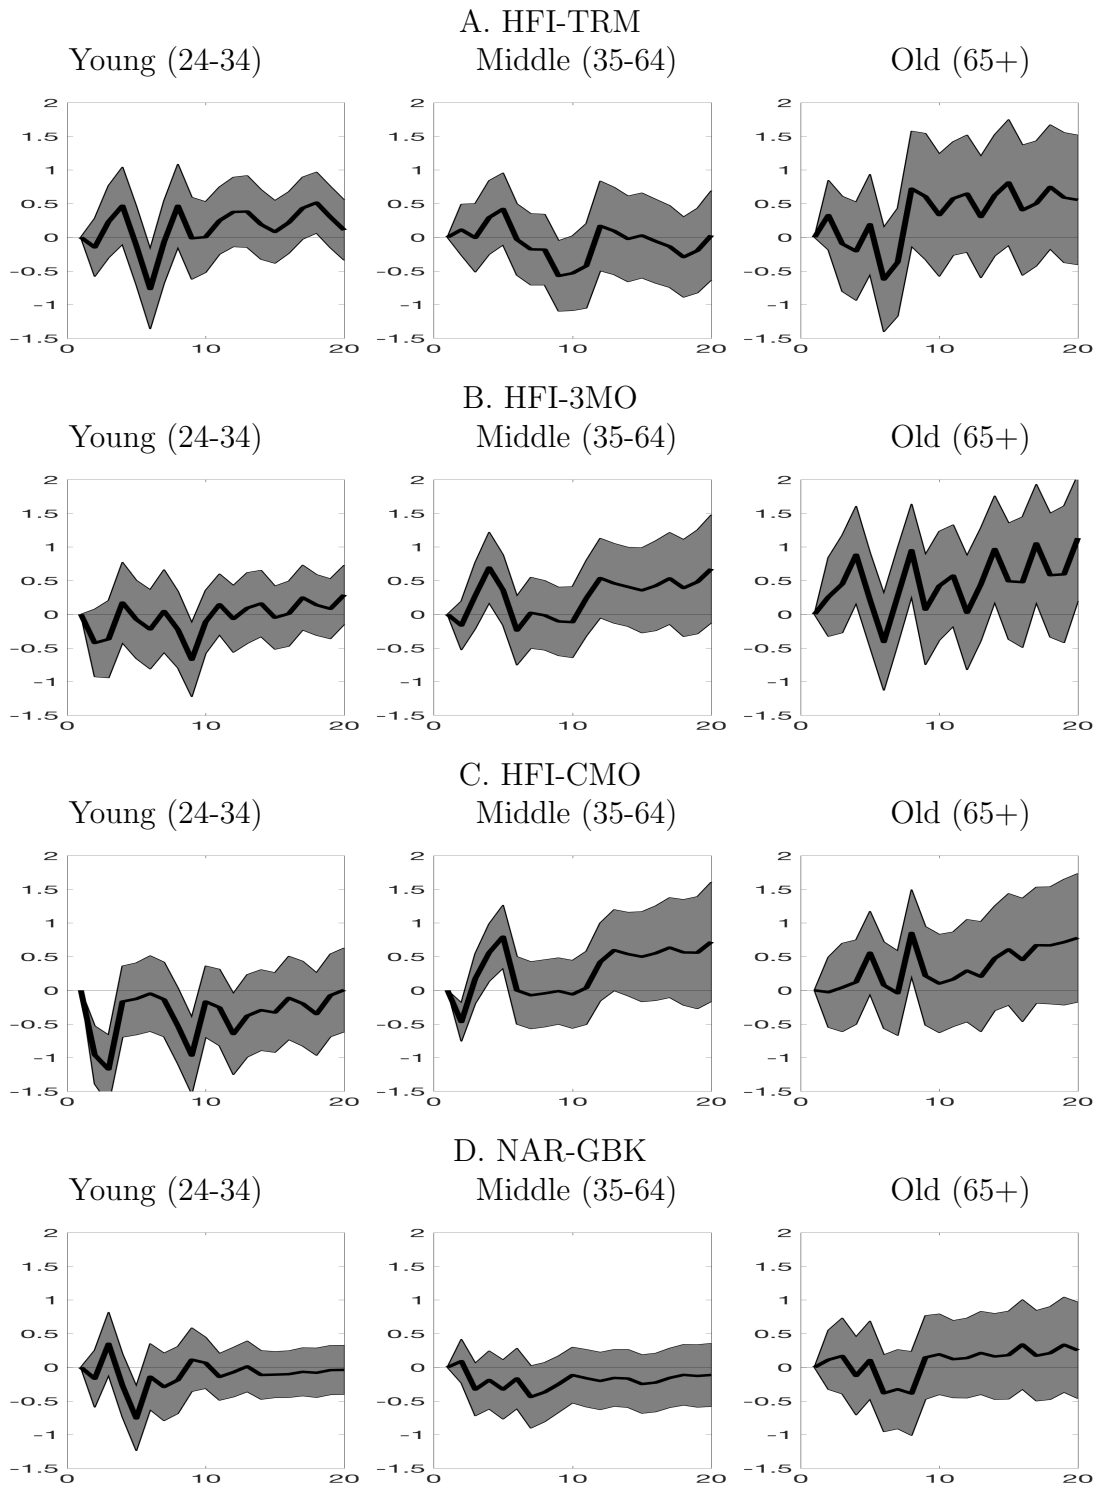
The consumption data studied to this point includes durables, whose purchases may be debt financed. Researchers and policy makers may also be interested in understanding patterns of non-durable consumption. To examine this, and to verify that our results are not driven entirely by durable expenditures, 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).<sup>11</sup>

Figure 7 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. In general, the size ordering of responses across all four shocks is  $\text{old} > \text{middle} \simeq \text{young}$ .

---

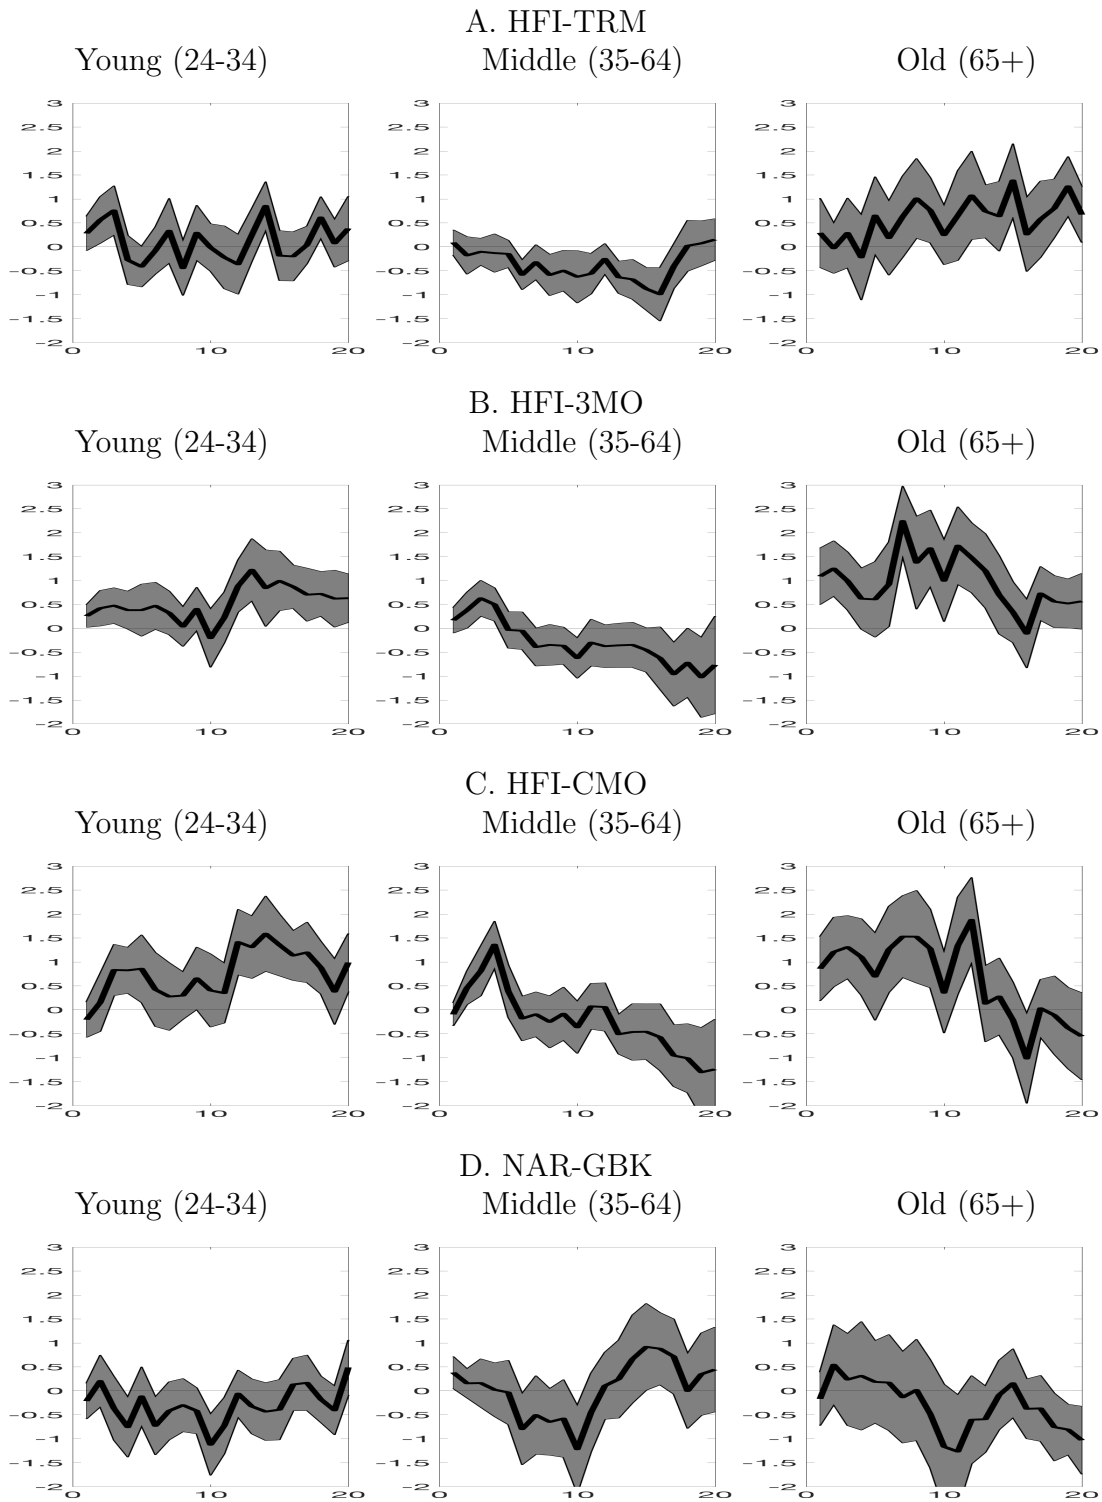
<sup>11</sup>Real non-durable consumption expenditures is the sum of consumption components 1-13 given in the Appendix A Table A-1.

Figure 7: Structural VAR – Cumulated Non-Durable Consumption Growth Impulse Response 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  $\pm$  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 8: Local Projections – Cumulated Non-Durable Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



Notes: This figure plots  $-\beta_h$  from Equation (4). Shaded areas are  $\pm$  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.

Figure 9: Structural VAR (SVAR) and Local Projections (LP) – Cumulated Non-Durable Consumption Growth Impulse Response for Old Minus Young and Old Minus Middle to Expansionary Monetary Policy Shock

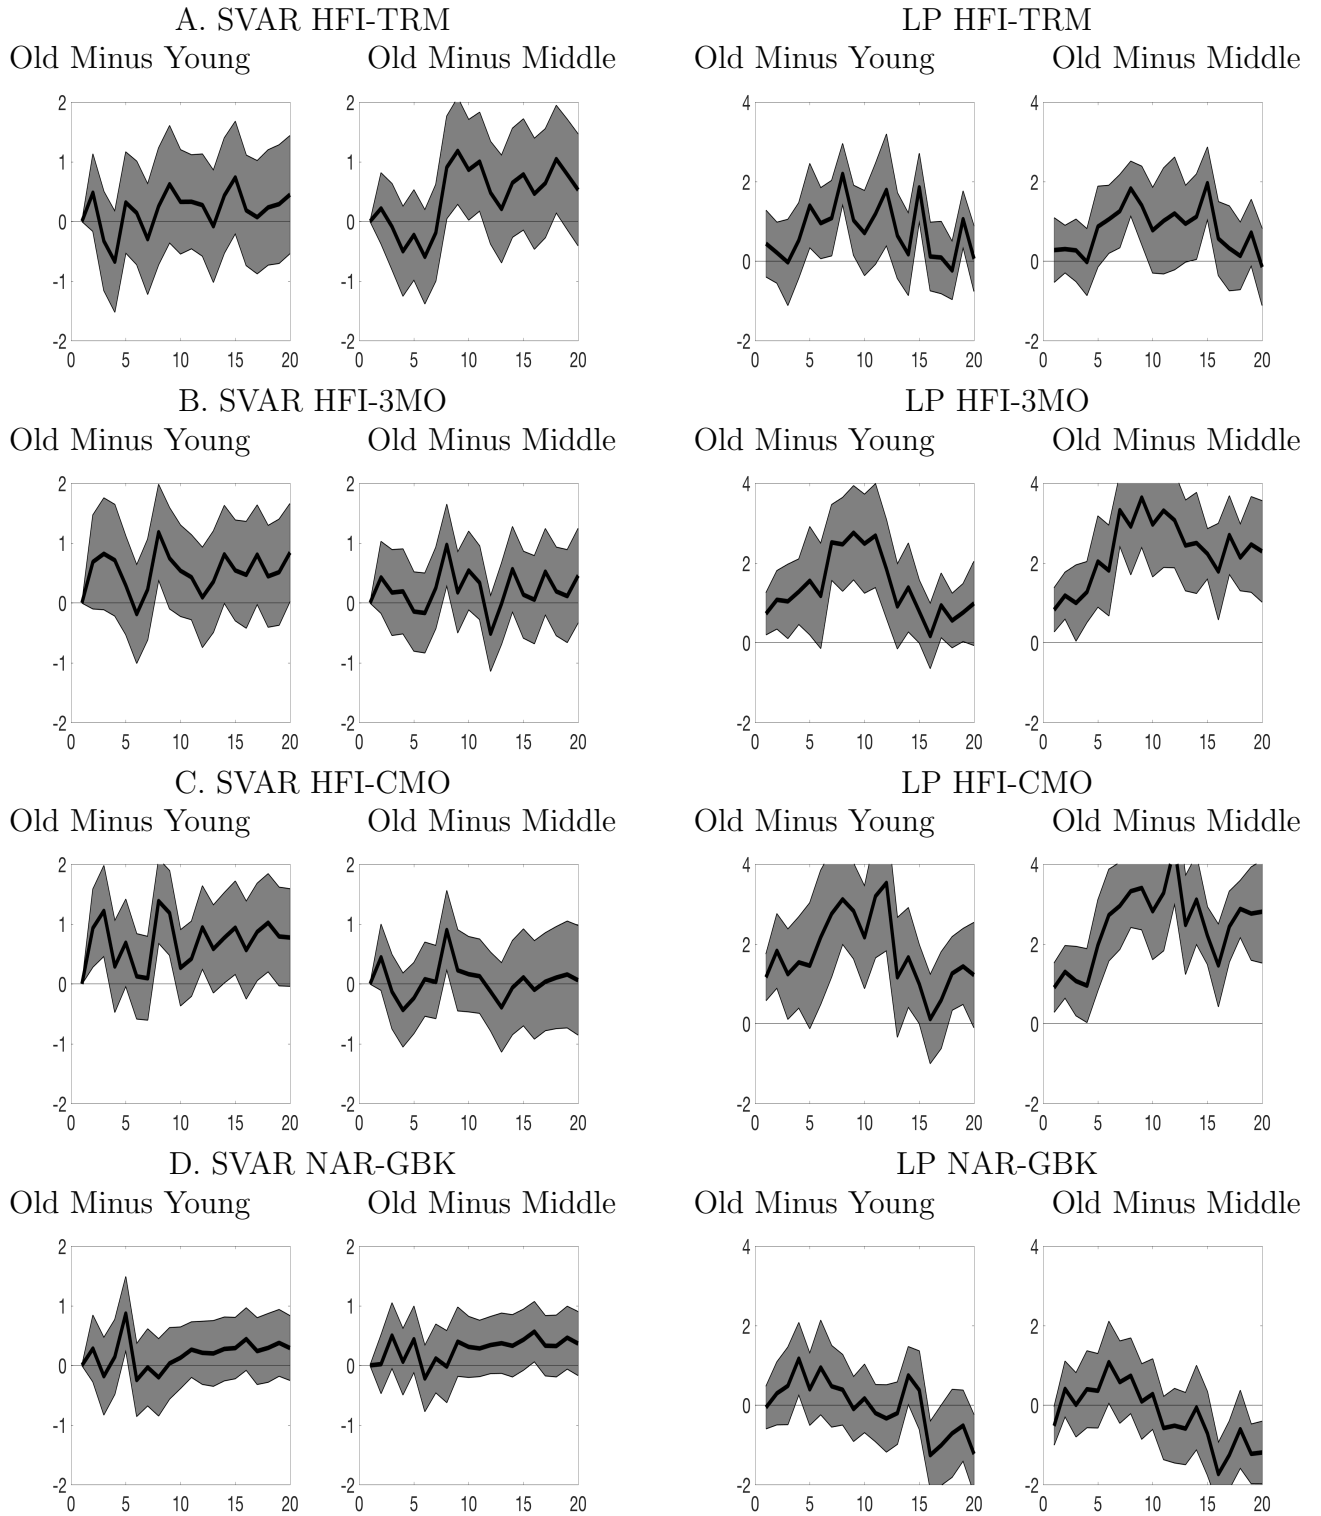


Figure 8 shows the corresponding local projection results. Again, the responses to all four shocks are similar. With the possible exception of the NAR-GBK shock, the response among the old remains the greatest.

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. Figure 9 displays the results in terms of the difference in consumption responses between the old households and the younger groups. Based on these findings, 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.

*Additional Robustness Checks.* Appendix B reports the following additional robustness checks:

1. Classification of households into six age groups.
2. Response of durable goods expenditure.
3. Response of consumption less housing expenditure.
4. Application of CEX weights in age group consumption aggregation.
5. The VAR in levels.
6. The VAR with  $k = 6$  lags.
7. Local projections without lagged consumption or the real federal funds rate.

Throughout the many specifications using structural VARs and local projections and across the four shocks, the message remains that aggregate consumption of the old households is the most responsive to monetary policy shocks. We next document key income and wealth facts by age group.

## 4 Income, Wealth, and Portfolio Composition by Age

Above, we report that consumption of the old is most responsive to monetary policy shocks. This section examines the relation among income, wealth, portfolio composition, and age using data from the 1989, 1998, and 2007 waves of the U.S. Survey of Consumer Finances (SCF) to explore potential mechanisms driving our empirical findings. The span of the sample approximately overlaps the time coverage of our CEX sample.<sup>12</sup>

---


<sup>12</sup>The 1989 SCF survey is the first that allows us to identify holdings in stock mutual funds and annuities.



The picture that emerges from this analysis 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. This points to heterogeneity in wealth and labor-supply across age groups as a potential mechanism driving the variation in consumption response. Old households adjust to the monetary policy shock induced wealth shock primarily by adjusting consumption whereas younger households can adjust both consumption and labor supply. Because old households have higher net wealth, they have higher exposure to the wealth shock than younger households. Auclert (2019) also makes this point, theorizing that households facing greater unhedged interest rate exposure respond more to monetary policy shocks. These differences combined with differences in marginal propensities to consume are his mechanism in propagating monetary policy shocks.

*Labor Income across Household Age.* Figure 10 shows the median and mean wage income as a share of total income by 5 year age groups. We are looking to see where in the life-cycle labor income is replaced by other sources. 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 pays for consumption using other sources.

Figure 10: 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 versus short-term assets across households of different ages. For each sample year and age group, we construct three measures of net asset positions.

1. *Net Worth* is total assets as stated by the SCF minus total debts.
2. *Net Long Term 1*, is the sum of net property equity (value of properties, including own residence, less outstanding debt on the properties), stock holdings, stock mutual funds, and annuities - a measure of long term, interest sensitive assets.
3. *Net Long Term 2* is *Net Long Term 1* plus non-stock mutual funds (but not money market funds) plus directly held bonds of all types.

Table 2 reports the age-group median *Net Worth* relative to the overall median, from which we see net wealth increasing with age. Since older age groups hold more wealth, they have higher exposure to valuation changes induced by monetary policy shocks.

Table 2: Within-Age Group Median Net Worth Relative to Aggregate Median Net Worth

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. The table reports 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 reports *Net Long Term 1* and *Net Long Term 2* as a fraction of *Net Worth* (the median of the ratio across households) by age group and year. Households with higher fractions of these measures have higher proportions of net wealth composed of long-term assets, and therefore higher exposure to interest rate fluctuations. The table shows that the share of long-term assets in net worth increases consistently with age. The largest component of net asset holdings for each age group is in property equity.<sup>13</sup> These data show that older households not only have higher net wealth but also that they hold more interest-rate sensitive assets.

<sup>13</sup>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. These results are available upon request.


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

Age Group	$\frac{\text{Net Long Term 1}}{\text{Net Worth}}$			$\frac{\text{Net Long Term 2}}{\text{Net Worth}}$		
	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: The table reports *Net Long Term 1* and *Net Long Term 2* as a fraction of *Net Worth* (the median of the ratio across households) by age group and year. *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.

We would like to estimate the VARs using consumption stratified by age and wealth (or a measure of interest rate exposure). Unfortunately, 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. We recognize that some people may object because doing so confounds income and wealth effects, particularly for working-aged households. Additionally, the income stratification can misclassify those hand-to-mouth working households who are also high wealth (Kaplan et al. (2014)). Nevertheless, due to lack of access to time-series data on wealth that can be linked to the consumption data, we proceed with income as a proxy for wealth.

Figure 11: Log Net Wealth and Log Income in 1989, 1998, and 2007



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

Some motivation for doing so is presented in Figure 11, which plots log net worth against log

income from the SCF in 1989, 1998, and 2007.<sup>14</sup> The size of each bubble is the representative SCF weight for that household. The vertical and horizontal axes are 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.<sup>15</sup>

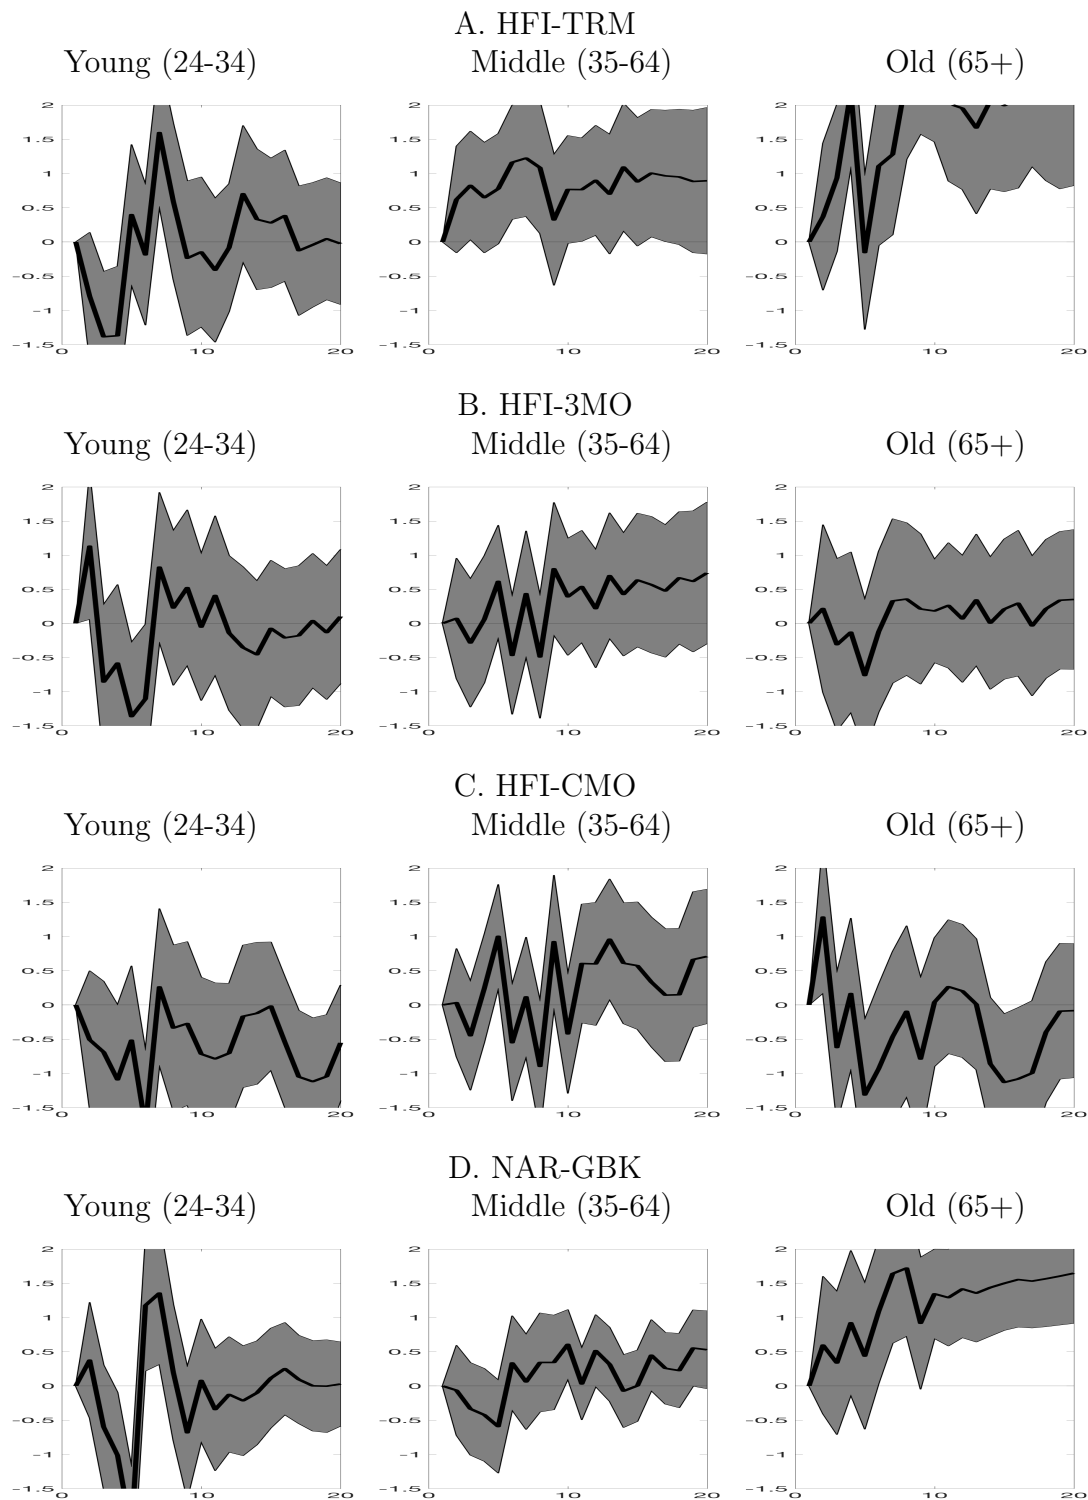
Accordingly, to implement this investigation, we split households in each age group into high and low income, based on whether their income is above or below the within-age group median income. Figure 12 shows the cumulated high-minus-low income household consumption impulse responses from the structural VAR by age group to an expansionary monetary shock. We see definitive differences between high and low income consumption responses to the HFI-TRM and NAR-GBK shocks for the old and a bit less so for the middle-aged. There is less separation for the old and middle in response to the other two shocks. The young display little separation in response to all of the shocks.

---

<sup>14</sup>We drop negative net worth observations. Due to the extreme wealth observations in the tails, we use logs to visualize the relationships.

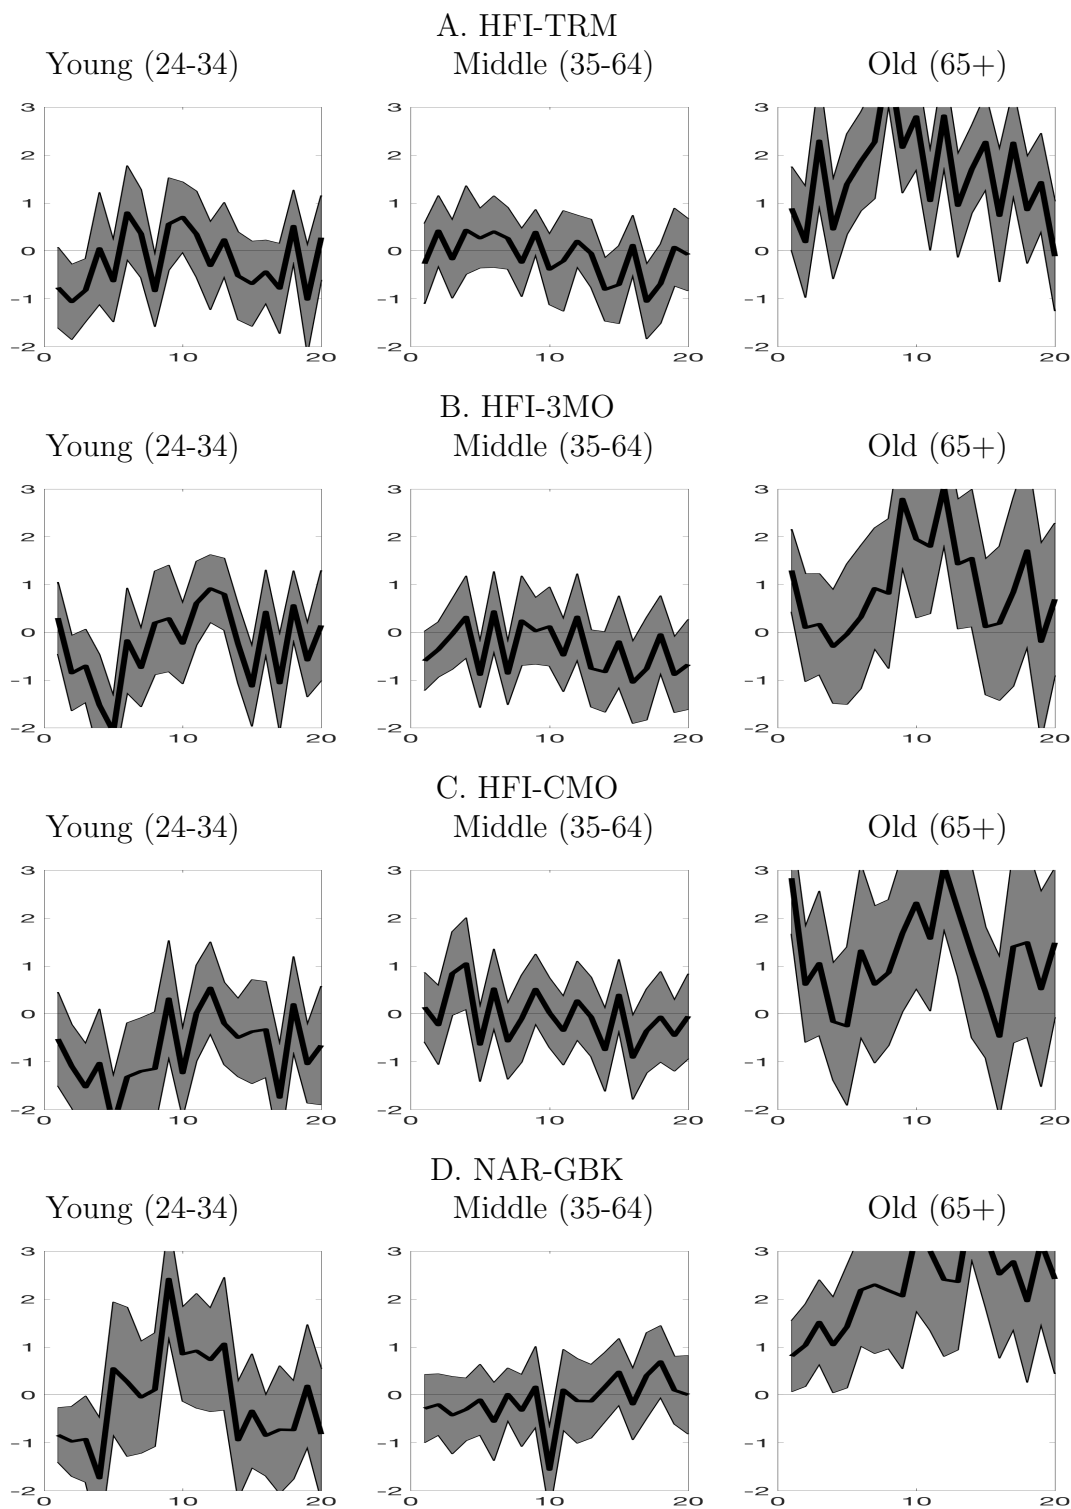
<sup>15</sup>In line with these estimates, Ríos-Rull and Kuhn (2016) find that the correlation between wealth and income is 0.58 in the 2013 SCF (see Table 19).

Figure 12: Structural VAR – Cumulated High Minus Low Income Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



Notes: High (low) income is average per capita consumption above (below) median income 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  $\pm$  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 13: Local Projections – Cumulated High Minus Low Income Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



Notes: High (low) income is average per capita consumption in the top (bottom) income decile by age group. Shaded areas are  $\pm$  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.

Figure 13 shows the corresponding high-minus-low income household consumption responses from local projections. The local projections yield more systematic separations between high and low income households for the old across all four shocks, and little separation for middle and young households.

These IRFs may be consistent with a number of stories, but they are also consistent with the idea that the young and middle-aged are relatively low net-wealth households whose consumption is paid mostly with labor income. Hence, after an interest rate shock, consumption for younger high and low income households react similarly. Conversely, income for old households is generated primarily from asset payoffs, such as interest and dividends, so the high (low) income old are also high (low) net wealth households. The rich and poor spending patterns differ from each other because the high (low) wealth households have more (less) valuation exposure to interest rates.

## 5 Divergence between Aggregate and Average Responses

We have presented evidence using structural VARs and local projections that the *aggregate* within age group consumption response is highest for older households. Wong (2018), on the other hand, using micro-level regressions, reports that the highest *average* household consumption response to monetary policy shocks is from young households. Both sets of results (aggregate and average responses) can be true if the consumption distribution and the distribution of the strength of response to monetary policy shocks are not symmetric.

Here is a simple example to illustrate how the two measures can diverge. Let there be two agents. Agent 1 consumes 99 and Agent 2 consumes 1. Aggregate consumption is 100 and agents are the same age. Following a shock, Agent 1 consumes 94 and Agent 2 consumes 2. The *aggregate* response is a 4 percent decline from 100 to 96. This is what we are after - the aggregate response within an age group. The *average* response in growth rates, however, is about a 47 percent increase (-5 percent for Agent 1, and +100 percent for Agent 2).

Subsection 5.1 illustrates this point with a small Monte Carlo experiment. Subsection 5.2 compares estimated average and aggregate consumption responses. We emphasize though, that the *average* effect is not central to our analysis, as we are interested in understanding the effect of monetary policy on the aggregate economy which naturally means studying the *aggregate* responses.

## 5.1 Monte Carlo Experiment

Log consumption for household  $i$  evolves according to

$$\ln(c_{i,t}) = \ln(c_{i,t-1}) + \mu_i + \gamma_i r_t + \epsilon_{i,t}, \quad (5)$$

where  $\epsilon_{i,t} \stackrel{iid}{\sim} N(0, \sigma_\epsilon^2)$ ,  $r_t$  is the real federal funds rate, and  $\mu_i$  is the household-specific consumption growth rate. The consumption response to the interest rate,  $\gamma_i$ , also varies by household.

The relationships amongst initial consumption,  $c_{i,0}$ ,  $\gamma_i$ , and  $\mu_i$  are governed as follows. Let  $X_i$  be a 3 by 1 zero-mean normal random vector,  $X_i \sim N(\omega, \Sigma)$ . Then  $\mu_i$  is normally distributed and  $\gamma_i$  and  $c_{i,0}$  are log-normally distributed,  $\mu_i = X_{i,1}$ ,  $\gamma_i = -e^{X_{i,2}}$ , and  $c_{i,0} = e^{X_{i,3}}$ , where the dependence among the three variables is specified in  $\Sigma$ . This ensures that realizations of  $\gamma_i$  are all non-positive and for  $c_{i,0}$ , all non-negative.

The size of the shock set for equation (5) was drawn from the data. The mean vector  $\omega$  and covariance matrix  $\Sigma$  of the normal distribution were set to match characteristics we found in the data. Values were set to directly target the average household consumption growth,  $\mu$ , and the mean and variance of log consumption,  $\ln(c_{i,0})$ , in 1986.<sup>16</sup>

Our analysis of high and low income households revealed additional patterns in the data, which we seek to capture in the simulations. Consumption growth for low-income young tends to be higher than for the high-income young, while the opposite is true of the old. Hence, we specify a small negative correlation between young-household consumption growth  $\mu$  and initial consumption, where high consumption presumably means high income, and a slightly larger but positive correlation for old households. We also found, in Section 4, that consumption of high-income old is more responsive than low-income old. The data is less definitive for the young, but we make a tentative argument that consumption of the high-income young is less responsive to monetary policy shocks than for the low-income young. Thus, we specify a positive correlation between initial consumption and responsiveness ( $\gamma$  is always negative) for the young, and a negative correlation for the old. Table 4 shows the implied moments for household specific growth, responsiveness to interest rate shocks, and initial consumption.

---

<sup>16</sup>This is the first year of usable observations in our VAR and Local Projections. The first two years are not used because of our use of 8 lags.



Table 4: Monte Carlo Experiment Data Generating Process Settings

	Young			Old		
	$\mu$	$\gamma$	$c_0$	$\mu$	$\gamma$	$c_0$
Mean	0.0013	-0.0480	2017	0.0023	-0.02726	1835
Std.Dev	0.001	0.0810	2079	0.00049	0.11952	2095
Skewness	0	-4.9	3.2	0.057	-9.36	3.9
Kurtosis	3	39	19.5	2.887	108	27
$\sigma_\epsilon$	0.034			0.041		
<u>Correlations</u>						
$\mu$		-0.05	-0.15		0.07	0.22
$\gamma$			0.33			-0.49

Notes:  $\gamma$  and  $c_0$  are log-normally distributed.  $\mu$  is normally distributed. The table shows the implied moments for  $\mu$ ,  $\gamma$ , and  $c_0$ .

The experiments each are 5,000 replications of 725 households with time-series length of 128, to approximately match the time-series observations in our macro analysis.

We estimate a panel regression with fixed effects,

$$100\Delta \ln(c_{i,t}) = a_i + \beta_i r_t + \epsilon_{i,t}$$

5,000 times. Also, we aggregate consumption (in levels)  $C_t = \sum_{i=1}^{725} c_{i,t}$  and estimate the aggregate time-series version of the regression

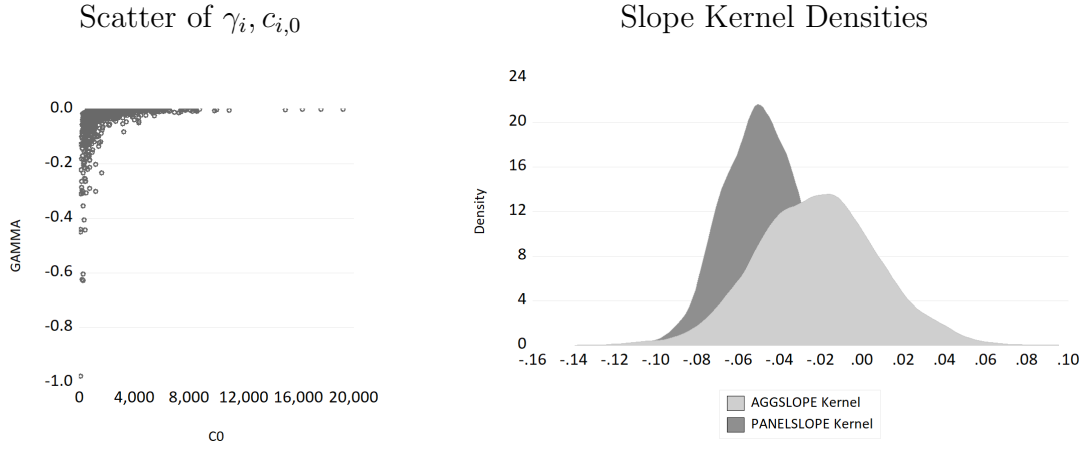
$$100\Delta \ln(C_t) = a + \beta r_t + \epsilon_t$$

5,000 times. Figure 14 shows the results.

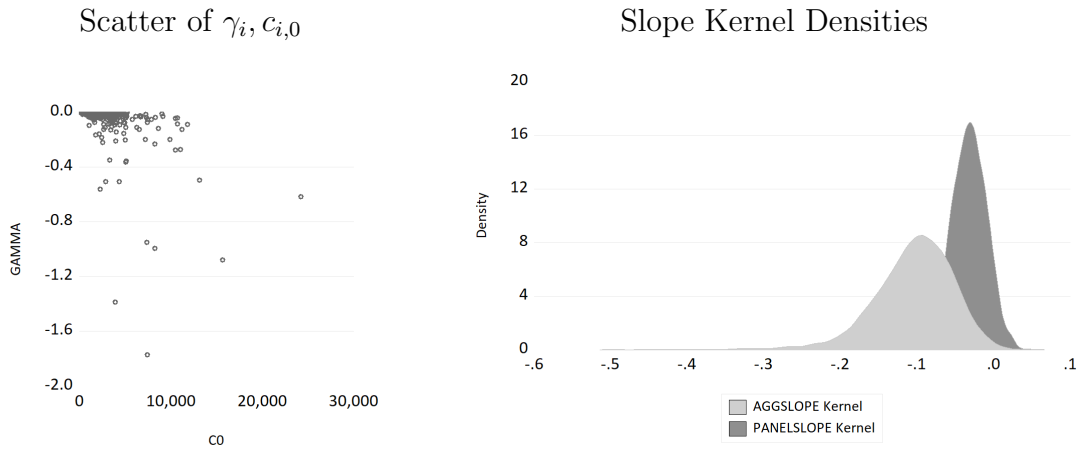
Panel A shows the scatter plot between  $\gamma_i$  and  $c_{i,0}$  and kernel densities of the fixed-effects estimate of the average household response and the aggregate time-series estimators for the experiment on the young. The mean of the aggregate time-series response is  $-0.02$ , whose magnitude is smaller than the  $-0.05$  mean of the average household response estimator. In Panel B, which shows results from the experiment for the old, we have a mean aggregate response of  $-0.10$  and a mean average response of  $-0.03$ . Thus, the monte carlo exercise demonstrates how the average (micro-level) impulse response can differ from the aggregate response.

Figure 14: Results of the Monte Carlo Experiment

A. Young Experiment



B. Old Experiment



## 5.2 Average and Aggregate Response Estimates

In this section, we compare our aggregate time-series estimates to panel regressions of the form

$$\Delta \ln(c_{i,t}) = b_0 + \sum_{k=1}^8 \beta_k s_{t-k} + X'_{i,t} \alpha + Z'_i \gamma + \lambda_{Q,t} + \epsilon_{i,t}, \quad (6)$$

used in studies of the average household response (e.g., Cloyne et al. (2018) and Wong (2018)). Appendix C provides the details about the construction of the micro data set.

Here,  $s_t$  is the monetary policy shock at time  $t$ . The vector  $X_{i,t}$  includes time varying household-level controls: changes in marital status of the household head (indicators capturing

unmarried to married, married to unmarried, or no change), changes in employment status of the household head (indicators capturing not employed to employed, employed to not employed, or no change), changes in household size, household size, changes in the number of persons 18 or younger in the household, and the number of persons 18 or younger. The vector  $Z_i$  includes time-invariant household controls: a household specific fixed effect (to control for any omitted household factors) and the birth year of the household head (to control for cohort effects). A set of quarterly indicators to control for seasonality is captured by  $\lambda_{Q,t}$ .

Table 5: Average And Aggregate Response of Young and Old,  $\sum_{j=1}^8 \beta_j$

	HFI-TRM	HFI-3MO	HFI-CMO	NAR-GBK
A. Average Response				
Young	0.403 (0.759)	-0.392 (0.623)	-0.533 (0.658)	-0.166 (0.745)
Old	-0.302 (0.634)	-0.835 (0.470)	-0.600 (0.572)	-0.889 (0.415)
B. Aggregate Response				
Young	0.381 (4.767)	2.996 (2.666)	5.520 (2.580)	0.113 (3.639)
Old	-2.783 (4.359)	-4.978 (3.379)	-6.750 (3.782)	-2.306 (3.296)

We suppress reporting results for middle-aged households, since the interesting contrasts are between the old and the young. Panel A in Table 5 shows the cumulated responses over eight quarters from estimating equation (6). We report cumulated responses over eight quarters in the spirit of Wong (2018), but we note that our peak consumption responses in our structural VARs and local projections occur beyond two years. Household observations are weighted by the CEX supplied weights.<sup>17</sup> Point estimates show old household consumption expands in response to each of the four expansionary monetary shocks. Average young household consumption is estimated to contract with HFI-TRM shocks. The average young are estimated to respond positively to expansions from the other three shocks, but by less than the old.

<sup>17</sup>The results are qualitatively unchanged if the CEX weights are omitted. If we instead weight by household consumption, then the old become even more responsive relative to the young.

Panel B reports the estimation results from

$$\Delta \ln(C_t) = b_0 + \sum_{k=1}^8 \beta_k s_{t-k} + \lambda_{Q,t} + \epsilon_t, \quad (7)$$

where  $C_t$  is aggregate consumption of a particular age group (young, middle, or old). Note that  $C_t$  is not the same measure of aggregate consumption used in the structural VARs and local projections. As explained in Appendix C, a large number of households must be dropped for this analysis, which are not dropped in the structural VAR and local projections analysis.

Equation 7 is estimated separately for each age group. The point estimates consistently show that the aggregate young household responses go in the ‘wrong’ direction following a monetary policy shock. The point estimates of the aggregate old household responses are consistently expansionary following expansionary shocks, as we found with the structural VARs and local projections. The aggregate response to monetary policy shocks is driven by older households.

As with the monte carlo experiment, the data reveals that the average response across households in an age group can differ from the aggregate response for that age group, in both size and sign. This empirical finding is consistent with our hypothesis that households living primarily off of their (high) wealth (who tend to be old households) are most responsive to interest rate changes. We next study this theory further by embedding it in a life-cycle model.

## 6 Wealth-Effects and Consumption Heterogeneity in a Life-Cycle Model

This section presents an overlapping generations model of finitely-lived households to illustrate how heterogeneity in net wealth, labor-supply choices, planning horizons, and discounting of the future might explain the observed consumption dynamics across age groups. The model agents save by accumulating long-term bonds both for retirement and to hedge against idiosyncratic income shocks. Preferences are given by Epstein and Zin (1989)–Weil (1989) recursive utility. Younger agents can supply labor while older agents are retired and survive on pension and asset payoffs. With these key ingredients, the model can replicate the higher consumption response among the old households and other features of the empirical impulse response functions reported

above. The model does not include liquidity or financial constraints.<sup>18</sup>

Once retired, model agents face an uncertain death and live a maximum of 86 years (344 quarters or periods in the model). People begin making economic decisions when they are 25 years old and enter economic life with no assets. 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 in debt. When people turn 65, they retire, face uncertain death, and live off of reduced pension income ( $S$ ) and accumulated assets. Retirees also have a bequest motive.<sup>19</sup> 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.

## 6.1 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 possible quarter of life.

The characteristics of the income process differ across each of the age groups. 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.

The transitory income shock  $u_{i,z,t}$ , is drawn from a mixture of a normal distribution and a

---

<sup>18</sup>Others (e.g., Parker et al. (2013)) have argued that such frictions matter for the consumption responses by age.

<sup>19</sup>We do not explicitly link cohorts; upon death, the bequests simply disappear and are not distributed to younger agents.

low-probability event of zero income for that quarter

$$u_{i,z,t} = \begin{cases} N(\mu_u, \sigma_u^2) & \text{with probability } (1-p) \\ -\infty & \text{with probability } p \end{cases} \quad (8)$$

where  $p$  is the probability of drawing zero income, and  $\mu_u = \frac{-\sigma_u^2}{2} - \ln(1-p)$ . 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

$$W_{i,z,t} = Y_{i,z,t} e^{u_{i,z,t}} \text{ for } z < z^* \quad (9)$$

and pension income for retired cohorts is

$$S_{i,z,t} = \begin{cases} Y_{i,z,t} & z = z^* \\ Y_{i,z,t} e^{u_{i,z,t}} & z > z^* \end{cases} . \quad (10)$$

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 movement along the age-earnings profile. The gross growth rate along this profile at cohort  $z$  is  $G_z$ . In retirement, both  $M_g$  and  $G_z$  become 1.

Let  $n_{i,z,t} \stackrel{iid}{\sim} N(\mu_n, \sigma_n^2)$  be the shock to permanent income  $Y_{i,z,t}$  and  $p_{rr}$  be the replacement rate on pension income. Then the life-cycle of permanent income evolves according to

$$Y_{i,z,t} = \begin{cases} Y_{i,z-1,t-1} M_g G_z e^{n_{i,z,t}} & z < z^* \\ p_{rr} Y_{i,z^*,t-1} & z = z^* \\ Y_{i,z-1,t-1} e^{n_{i,z,t}} & z > z^* \end{cases} . \quad (11)$$

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). We select data between 1986 and 2007 to align with the span of our CEX sample. We use the same definition of household income in the PSID as Blundell et al. (2008) and Storesletten et al. (2007). Our estimation method follows Choi et al. (2017), who build on Zeldes (1989), Carroll (1992), and Carroll (1997).

From the raw income data, we first remove the aggregate time trend, predictable life-cycle or occupation dependent fluctuations, and household fixed effects. The remaining variation is used to estimate the parameters  $(\sigma_n, \sigma_u, p)$  separately for young, middle, and old households. 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 6: 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		
$\sigma_u$	0.471	0.467	0.482		
$\sigma_n$	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,  $\sigma_u$  is the standard deviation of transitory income, and  $\sigma_n$  is the standard deviation of permanent income.

Table 6 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, we ‘forecast’ values for ages 61-64 with a cubic trend regression. As seen from the table, income

peaks at age 51 and macroeconomic income growth is virtually nil, with an annual growth rate of 60 basis points.

Panel C shows the remaining parameters  $(p, \sigma_u, \sigma_n)$ , 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, with  $\sigma_n = 0.14$ .

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). Overall, the shape of the resulting life-cycle income process is in line with other recent estimates in the literature (see Guvenen et al. (2015) and Guvenen et al. (2018), for example) based on alternative methodologies.

## 6.2 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 cohort  $z$ , at time  $t$ . Labor supply is  $L_{i,z,t}$  and, normalizing the time endowment to 1, leisure is  $(1 - L_{i,z,t})$ .

Working age household  $z < z^*$  utility is,

$$V_{i,z,t} = \left\{ (1 - \beta) (C_{i,z,t}^\nu (1 - L_{i,z,t})^{1-\nu})^{(1-\rho)} + \beta \left[ (E_t [V_{i,z+1,t+1}^{1-\gamma}])^{\frac{(1-\rho)}{(1-\gamma)}} \right] \right\}^{\frac{1}{(1-\rho)}} \quad (12)$$

$\rho > 0$ ,  $\gamma > 0$ ,  $0 \leq \nu \leq 1$ , and  $0 < \beta < 1$ .  $\beta$  is the subjective discount factor.  $\rho^{-1}$  is the intertemporal elasticity of substitution.

Households face idiosyncratic income risk and live in an incomplete markets environment. Neither contingent claims nor insurance instruments are available.<sup>20</sup> The non-human asset is a non-state contingent long-term (consol) bond, that pays one unit of consumption each period forever. The intent is for the long-term bond to mimic 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) asso-

---

<sup>20</sup>Our estimation of income shocks in the data was net of all transfers and thus corresponds to the notion of uninsurable risk in the model.



ciated with housing. Additionally, our previous analysis of the SCF data revealed that older households hold many other long-term and interest-rate-sensitive financial instruments.

A working-aged household can borrow or lend by going short or long the bond.<sup>21</sup> The net number of bonds held by the household is  $A_{i,z,t}$ . Upon retirement, households face a possibility of death and must have non-negative assets in retirement to ensure that they do not die in 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. 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 consumption. Their budget constraints are

$$P_t^a A_{i,z+1,t+1} = A_{i,z,t} + P_t^a A_{i,z,t} + L_{i,z,t} W_{i,z,t} - C_{i,z,t}, \quad (13)$$

which can be written in a more familiar form,

$$A_{i,z+1,t+1} = A_{i,z,t} + r_t (A_{i,z,t} + L_{i,z,t} W_{i,z,t} - C_{i,z,t}). \quad (14)$$

Retired households have a bequest motive, supply no labor and face an uncertain death where the cohort  $z$  specific probability of surviving to age  $z+1$  is  $\delta_{z,t}$ . Following Gomes and Michaelides (2005), we model the bequest motive of retirees as,

$$\frac{1}{1-\gamma} E_t \left( \left( \frac{1}{b} \frac{A_{i,z+1,t+1}}{r_t} \right)^\nu \right)^{1-\gamma}.$$

Hence, utility for retired households, aged  $z^* \leq z < Z = 248$ , is

$$V_{i,z,t} = \left\{ (1 - \beta \delta_{z,t}) (C_{i,z,t}^\nu)^{(1-\rho)} + \beta \left[ \delta_{z,t} E_t (V_{i,z+1,t+1}^{1-\gamma}) + \frac{(1 - \delta_{z,t})}{1-\gamma} E_t \left( \left( \frac{1}{b} \frac{A_{i,z+1,t+1}}{r_t} \right)^\nu \right)^{1-\gamma} \right]^{\frac{1-\rho}{1-\gamma}} \right\}^{\frac{1}{1-\rho}}, \quad (15)$$

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

---

<sup>21</sup>Cash flow effects related to interest rate changes are confined to rebalanced asset holdings. In contrast, Auclert (2019) emphasizes unhedged interest exposures, which arise both from maturing assets and liabilities, as well as short-term assets. Appendix E reports the results when the saving instrument is instead a one-period bond.

$z = Z = 248$ , and  $\delta_{248,t} = 0$ , so utility is

$$V_{i,Z,t} = \left\{ (C_{i,z,t}^\nu)^{(1-\rho)} + \beta \left[ \frac{1}{1-\gamma} E_t \left( \left( \frac{1}{b} \frac{A_{i,z+1,t+1}}{r_t} \right)^\nu \right)^{1-\gamma} \right]^{\frac{1-\rho}{1-\gamma}} \right\}^{\frac{1}{1-\rho}}, \quad (16)$$

Retired households face budget constraints

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

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

### 6.3 Solution and Parameterization

To solve the model, we discretize the state space and obtain policy functions for the stationary model where variables are normalized by permanent income. 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 D 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). We obtain the long-term interest rate from the short rate using the expectations theory of the term structure with two modifications. 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. With 5 states for the interest rate, the long-term rate can take values of 3.08, 3.27, 3.44, 3.61, and 3.78 percent per annum.

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  $\beta = 0.9962$ ,  $\gamma = 12$ ,  $\nu = 0.5$ , and  $\rho = 1.4^{-1}$ . This gives an annualized rate of time preference of 1.54% and an intertemporal elasticity of substitution of 1.4 (which is the value estimated in the working paper version of Choi et al. (2017) for American households).<sup>22</sup>

---

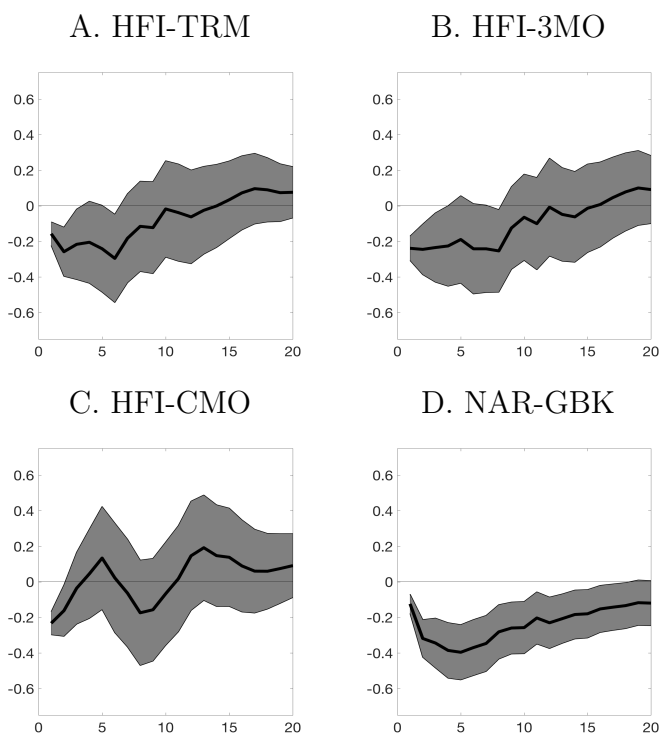
<sup>22</sup>With infinitely lived agents, a value of  $\nu = 0.36$  typically gives a steady state choice of time worked at 1/3 of the time endowment. We set  $\nu$  at a slightly higher value.

## 6.4 Model Impulse Responses

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

Empirically, the transmission mechanism runs from the monetary policy shock to the interest rate then to consumption decisions. In the model, we want the interest rate dynamics driving the model's impulse response to look as it does in the data. The way it looks in the data is shown in Figure 15, which displays the impulse responses of the real federal funds rate to the four monetary policy shocks from the structural VAR in Section 3.1.<sup>23</sup>

Figure 15: Structural VAR – Real Federal Funds Rate Impulse Response to Expansionary Monetary Policy Shock



Note: Results are from the structural VAR estimated in Section 3.1.

As can be seen, the shocks (except perhaps in the response to HFI-CMO) have persistent

---

<sup>23</sup>Romer and Romer (2004) and Coibion et al. (2017) among others, also find a high degree of persistence in response to monetary policy shocks.

effects on the real federal funds rate. To conform to the empirics, we model the expansionary interest rate shock as a persistent decline. For four periods before the shock, the long-term interest is set at its mean value (3.44%). At the time of the shock, it declines to its lowest value (3.08%) for 7 quarters, then rises to the next lowest value (3.27%) for the next 6 quarters, before resuming its random evolution.


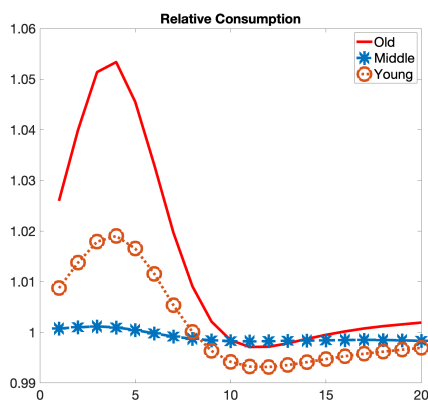
We simulate un-normalized responses of log consumption, log labor supply, asset quantities, and asset values for each individual to the negative interest rate shock.<sup>24</sup> We then take the mean within each age cohort, and then the mean within each of the three broad age groups.

Figure 16 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. If the shock occurs at  $t^*$ , 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 response for the young is also positive, though slightly smaller. The consumption response for the middle-aged group is small and nearly zero. Overall, the life-cycle model can replicate the main qualitative features of the consumption responses estimated from the CEX data.


---

<sup>24</sup>Assets are not logged since young and middle households can borrow, resulting in negative values for assets.

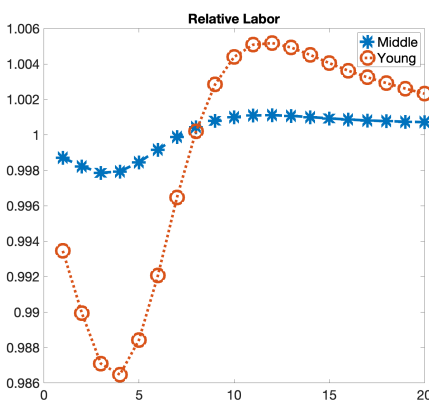
Figure 16: Relative Responses to Negative Interest Rate Shock  
 A. Log Consumption                      B. Number of Assets



C. Value of Assets



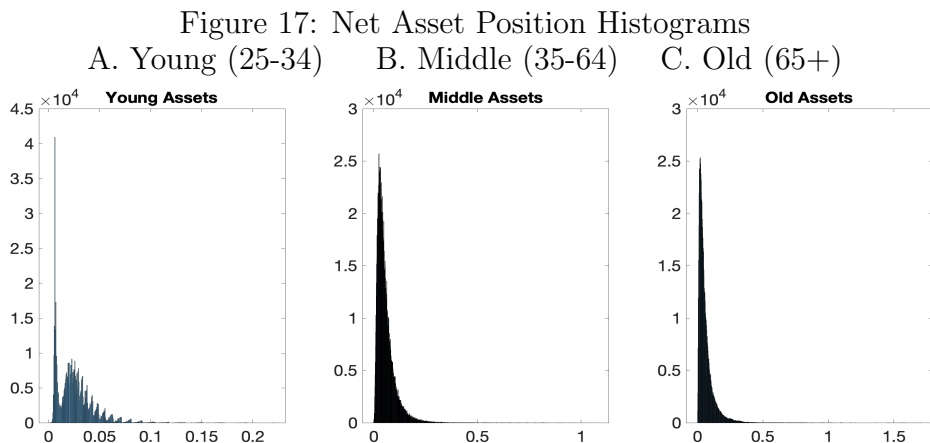
D. Log Labor



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. The persistence of the consumption response is induced by the persistent decline in the interest rate. Even though the young draw down relatively more assets than the old (Figure 16, Panel B), the old hold far more assets. The relative response patterns of asset values (Figure 16, Panel C) held by the different age groups are roughly similar, but the young get the smallest capital gains. Both the young and middle take more leisure and reduce labor in the first eight quarters following the shock, with a larger response by young households (Figure 16, Panel D). The interest rate cut causes both consumption and leisure for young households to increase, whereas the aggregate response of middle-aged households is more muted.

Figure 17 shows the histograms of the asset holdings for the three age groups. The distributions 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. Because the intertemporal elasticity of substitution is greater than 1, the young and middle-aged maintain positive net worth.




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 age-net worth pattern generated by the model mimics the pattern found in the SCF data. The young in the model hold few assets relative to the large number of assets held by the old. In the model, most households accelerate their asset holdings about ten years before they retire. This results in a mixed and muted impact on the middle group, similar to what we estimated in the structural VARs and local projections.

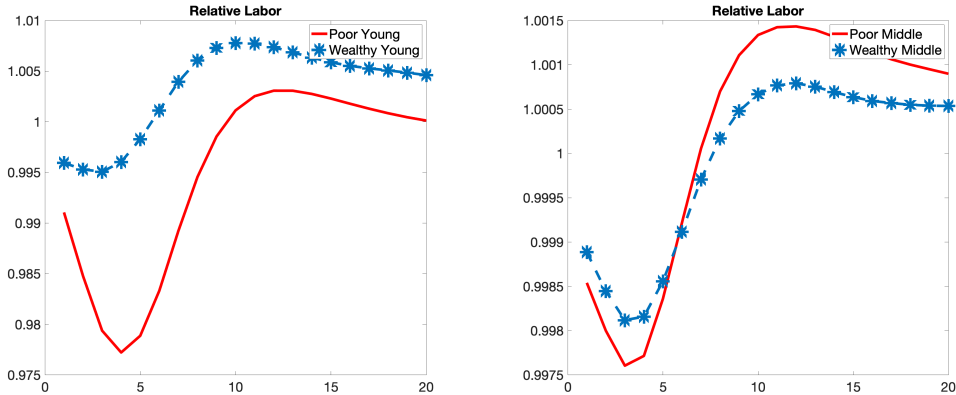
In Figure 18 we divide households into ‘wealthy’ (above median wealth) and ‘poor’ (below median wealth) for each age group, as we did with income in the empirical analysis of Section 4. The figure shows consumption of the wealthy old to be more responsive to interest rate shocks than consumption of the poor old. There is little difference between middle poor and wealthy (note the scale). In contrast, the poor young increase consumption by more than wealthy young, which is broadly consistent with the response patterns obtained in the structural VARs and local projections. Also, the poor young reduce labor by more than the wealthy young (Figure 19). This is due to the nonseparability between leisure and consumption in utility in which higher consumption is associated with higher leisure.

Figure 18: Relative Log Consumption Responses by Wealthy and Poor  
 A. Young (25-34)                      B. Middle (35-64)                      C. Old (65+)



Notes: The figure shows simulated un-normalized relative responses of mean log consumption across age groups for the wealthy and poor to a  $-0.36$  percent decline in the long-term bond rate. If the shock occurs at time  $t^*$ , the relative response for consumption is  $\ln(C_t/C_{t^*})$ . Wealthy households are those with above median asset holdings. The horizontal axis indicates the number of quarters for up to five years after the shock.

Figure 19: Relative Log Labor Responses by Wealthy and Poor  
 A. Young (25-34)                      B. Middle (35-64)



Notes: The figure shows simulated un-normalized relative responses of mean log labor across age groups for the wealthy and poor to a  $-0.36$  percent decline in the long-term bond rate. If the shock occurs at time  $t^*$ , the relative response for labor is  $\ln(L_t/L_{t^*})$ . Wealthy households are those with above median asset holdings. The horizontal axis indicates the number of quarters for up to five years after the shock.

To summarize, the model qualitatively explains key features of the data. In response to an exogenous interest rate cut, both the data and the model exhibit i) consumption by the old increases by more than the middle and young, ii) consumption by the wealthy old increases by more than the poor old, iii) not much of a difference between consumption by the wealthy and poor middle, iv) consumption of the young poor is higher than that of the wealthy poor.

We close this section by mentioning that the recursive utility structure combined with a labor-leisure choice for the young and middle-aged seem to be necessary to get the old to be the most responsive. Section E of the appendix shows model impulse responses under alternative parameter settings.

## 7 Conclusion

The weight of the evidence presented across alternative monetary policy shocks, empirical methods, and consumption measures is that aggregate consumption of old households react more to monetary policy shocks than do middle and young households. 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 feature of the data, which is that consumption of old households is 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. (2019): “Monetary Policy and the Redistribution Channel,” *American Economic Review*, *forthcoming*.
- 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.
- CLOYNE, J., C. FERREIRA, AND P. SURICO (2018): “Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism,” *Review of Economic Studies*, *forthcoming*.
- 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*.

- DOEPKE, M. AND M. SCHNEIDER (2006): “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 114, 1069–1097.
- DYNAN (2009): “The Effects of Population Aging on the Relationship among Aggregate Consumption, Saving and Income,” *American Economic Review*.
- EICHENBAUM, M., S. REBELO, AND A. WONG (2018): “State Dependent Effects of Monetary Policy: the Refinancing Channel,” *Working Paper*.
- 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*.
- GOMES, F. AND A. MICHAELIDES (2005): “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *The Journal of Finance*, 60, 869–904.
- GORNEMANN, N., K. KUESTER, AND N. MAKOTO (2012): “Monetary Policy with Heterogeneous Agents,” *FRB Philadelphia WP No. 12-21*.
- GORODNICHENKO, Y. AND M. WEBER (2016): “Are Sticky Prices Costly? Evidence from the Stock Market,” *American Economic Review*, 106, 165–99.
- GUREN, A. M., A. MCKAY, E. NAKAMURA, AND S. JÒN (2018): “Housing Wealth Effects: The Long View,” *University of California, Berkeley, Mimeo*.
- (2019): “What Do We Learn From Cross-Sectional Empirical Estimates in Macroeconomics?” *University of California, Berkeley, Mimeo*.
- 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.

- GUVENEN, F., G. KAPLAN, J. SONG, AND J. WEIDNER (2018): “Lifetime Incomes in the United States Over Six Decades,” *NBER Working Paper No. 23371*.
- GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2015): “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?” *NBER Working Paper No. 20913*.
- 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.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): “The Wealthy Hand-to-Mouth,” *NBER Working Paper No. 20073*.
- KRUEGER, D. AND F. PERRI (2006): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *The Review of Economic Studies*, 73, 163–193.
- LEAHY, J. V. AND A. THAPAR (2019): “Demographic Effects on the Impact of Monetary Policy,” *University of Michigan, Mimeo*.
- 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.
- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 1687–1726.
- MIAN, A. AND A. SUFI (2009): “The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis,” *Quarterly Journal of Economics*, 124, 1449–1496.
- (2011): “House Prices, Home Equity-based Borrowing, and the US Household Leverage Crisis,” *American Economic Review*, 101, 2132–2156.
- NAKAMURA, E. AND J. STEINSSON (2018a): “Identification in Macroeconomics,” *Journal of Economic Perspectives*, 32, 59–86.

- (2018b): “High Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Columbia University, Mimeo*.
- NEWKEY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703.
- PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 103, 2530–53.
- 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.
- STERK AND TENREYRO (2018): “The Transmission of Monetary Policy through Redistributions and Durable Purchases,” *Journal of Monetary Economics*.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2007): “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” *Review of Economic Dynamics*, 10, 519–548.
- 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. (2018): “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 The Consumption Data

Our consumption data is taken from the quarterly Consumer Expenditure Survey Interview Sample from 1984Q1 to 2007Q4, compiled by the U.S. Bureau of Labor Statistics. The data are found in the Family Characteristics and Income (FAMILY) files.

Following Krueger and Perri (2006), we construct real consumption expenditures by deflating the household’s expenditures for each of the 19 categories listed in Table A–1 by its category-specific deflator. The table also lists the BLS CPI code. Total household consumption is the sum of expenditures on components 1-19.

Let  $c_{i,t}$  be total household consumption divided by the number of members in household  $i$  at time  $t$ . Let  $H_t$  be the total number of household observations in the group. Following Anderson et al. (2016), within-group aggregate consumption is

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

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 <sup>†</sup>
15	House equipment	HOUSEQCQ	Household furnishings and operations	SAH3
16	Vehicles	TRANSCQ-GASMOCQ -PUBTRACQ	Private Transportation	SAT1 <sup>‡</sup>
17	Other lodging	OTHLDCQ	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.

<sup>†</sup> Prior to 1998 this is SA6 (Entertainment). From 1998 on it is SAR (recreation).

<sup>‡</sup> We combine purchases and vehicle maintenance into Vehicles in the CEX category and use private transport in the CPI.

*Generating Quarterly Consumption:* The CEX is a rotating survey where respondents are interviewed up to 5 times. Respondents are interviewed once a quarter, but the interview can

occur in any of the 3 months within that quarter. The first interview collects information on the household's characteristics, but not it's consumption expenditures. Hence, each households has at most 4 usable observations. In the subsequent interviews, the household reports expenditures over the previous 3 months. There is a difference between the calendar quarter and the interview quarter, and we discuss here how we calculate calendar quarter observations, which is illustrated in Table A-2.<sup>25</sup>

The table shows 4 fictitious households, HH-1 through HH-4, each interviewed in a different month. HH-1 is interviewed in January about its October through December expenditures. HH-2 is interviewed in February about its November-January expenditures, and so forth.

To align the interview-quarter expenditures to calendar-quarters, we treat expenditures in a given month as representative of expenditures for the calendar-quarter. We illustrate in Table A-3. HH-1's interview provides information for the entire 4th quarter of 2000, so the interpretation is clear. HH-2 is reporting spending for 2 months in 2000Q4. We multiply that spending number by 3/2 and that becomes HH-2's 2000Q4 consumption. HH-2 reports spending for 1 month in 2001Q1. That spending number is multiplied by 3 and that becomes HH-2's 2001Q1 consumption. The adjustments for HH-3 and HH-4 follow analogously.<sup>26</sup>

Table A-2: Interview month and calendar-quarter

Calendar - Quarter	Month of Expenditure Recorded	Month of Interview			
		Jan. 2001 HH-1	Feb. 2001 HH-2	Mar. 2001 HH-3	Apr. 2001 HH-4
2000 Q4	Oct.	✓			
	Nov.	✓	✓		
	Dec.	✓	✓	✓	
2001 Q1	Jan.		✓	✓	✓
	Feb.			✓	✓
	Mar.				✓

Table A-3: Calendar-Quarter Consumption

	HH-1	HH-2	HH-3	HH-4
2000 Q4	$(\text{Oct.} + \text{Nov.} + \text{Dec.}) \times \frac{3}{3}$	$(\text{Nov.} + \text{Dec.}) \times \frac{3}{2}$	$\text{Dec.} \times \frac{3}{1}$	
2001 Q1		$\text{Jan.} \times \frac{3}{1}$	$(\text{Jan.} + \text{Feb.}) \times \frac{3}{2}$	$(\text{Jan.} + \text{Feb.} + \text{Mar.}) \times \frac{3}{3}$

<sup>25</sup>The rotating design and difference between calendar and interview quarter is discussed on p.22 of <https://www.bls.gov/cex/2015/csxintvw.pdf>.

<sup>26</sup>We have also considered an alternative way to calculate quarterly consumption - using weights on how many months the consumption observation is in the sample. The results were not sensitive to this alternative.

## B Additional Empirical Results

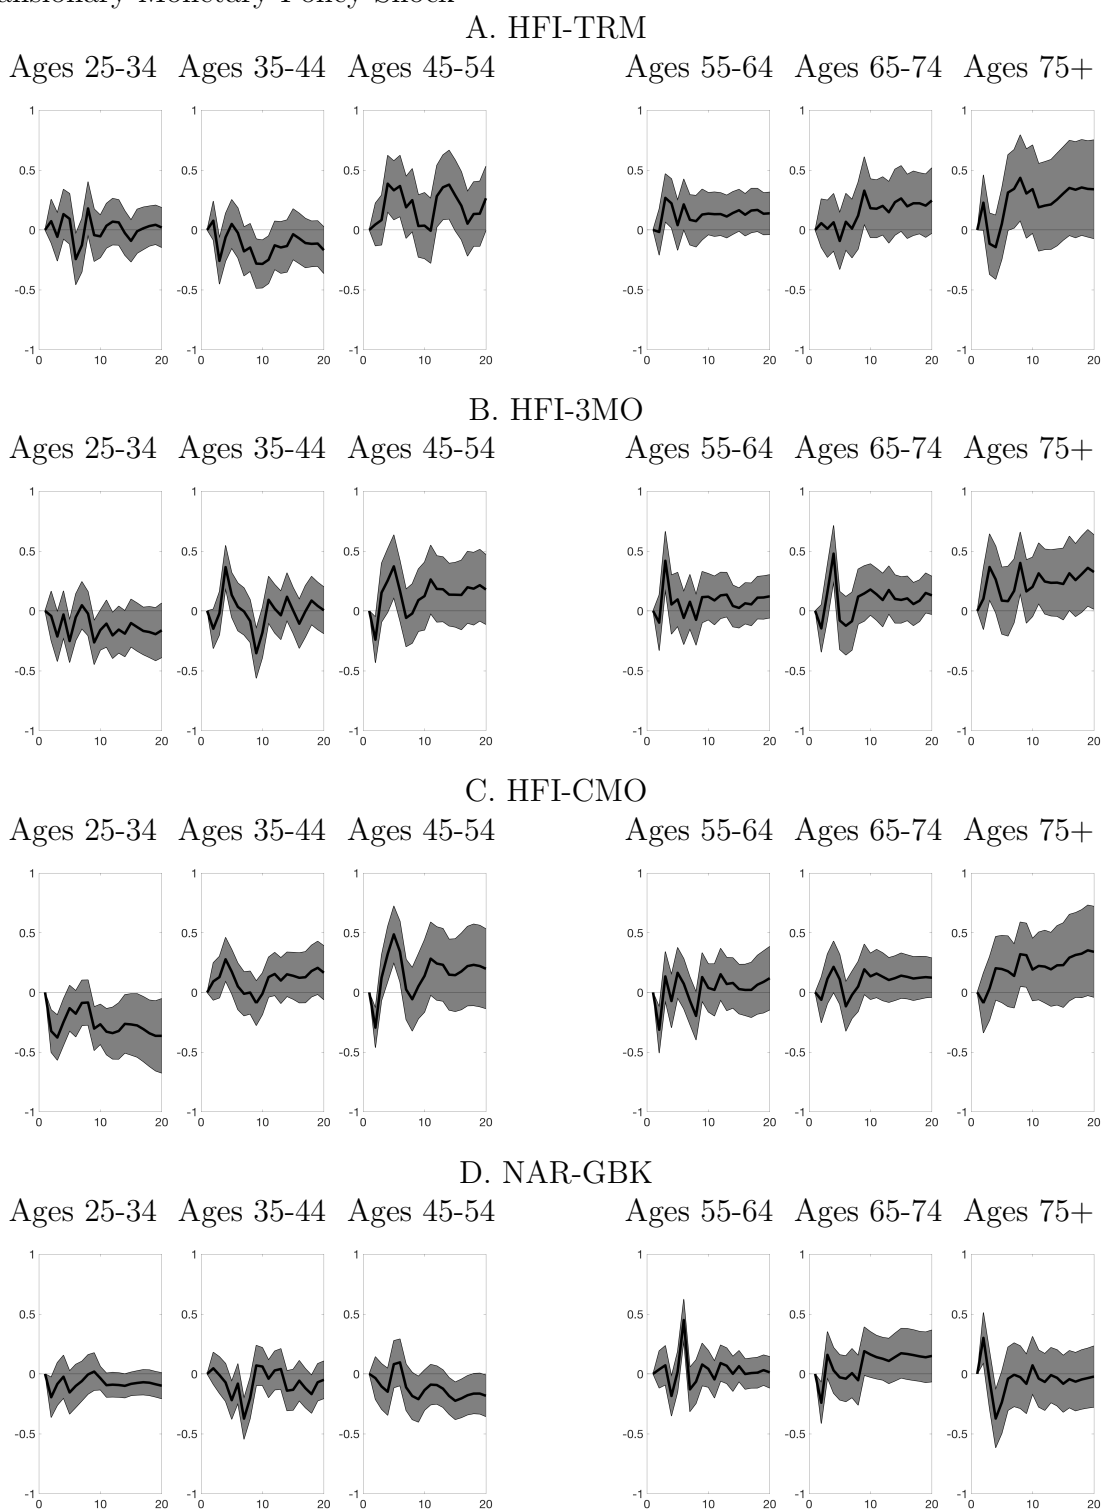
This section reports the following additional empirical VAR results.

1. Classification of households into six age groups (25-34, 35-44, 45-54, 55-64, 65-74, 75 and above).
2. Response of durable goods expenditure.
3. Response of consumption less housing expenditure.
4. Application of CEX weights in age group consumption aggregation.
5. The VAR in levels.
6. The VAR with  $k = 6$  lags
7. Local projections without controlling for lagged consumption or the real federal funds rate.

### 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 and B.2 report the consumption responses based on the structural VARs and local projections and one standard deviation expansionary HFI-TRM, HFI-3MO, HFI-CMO, and NAR-GBK monetary policy shocks. Once we split the sample into finer groups, we again see the old respond most to monetary policy shocks, although in a few instances the 65-74 age group responds most. The responses, in general, get progressively stronger as households age.

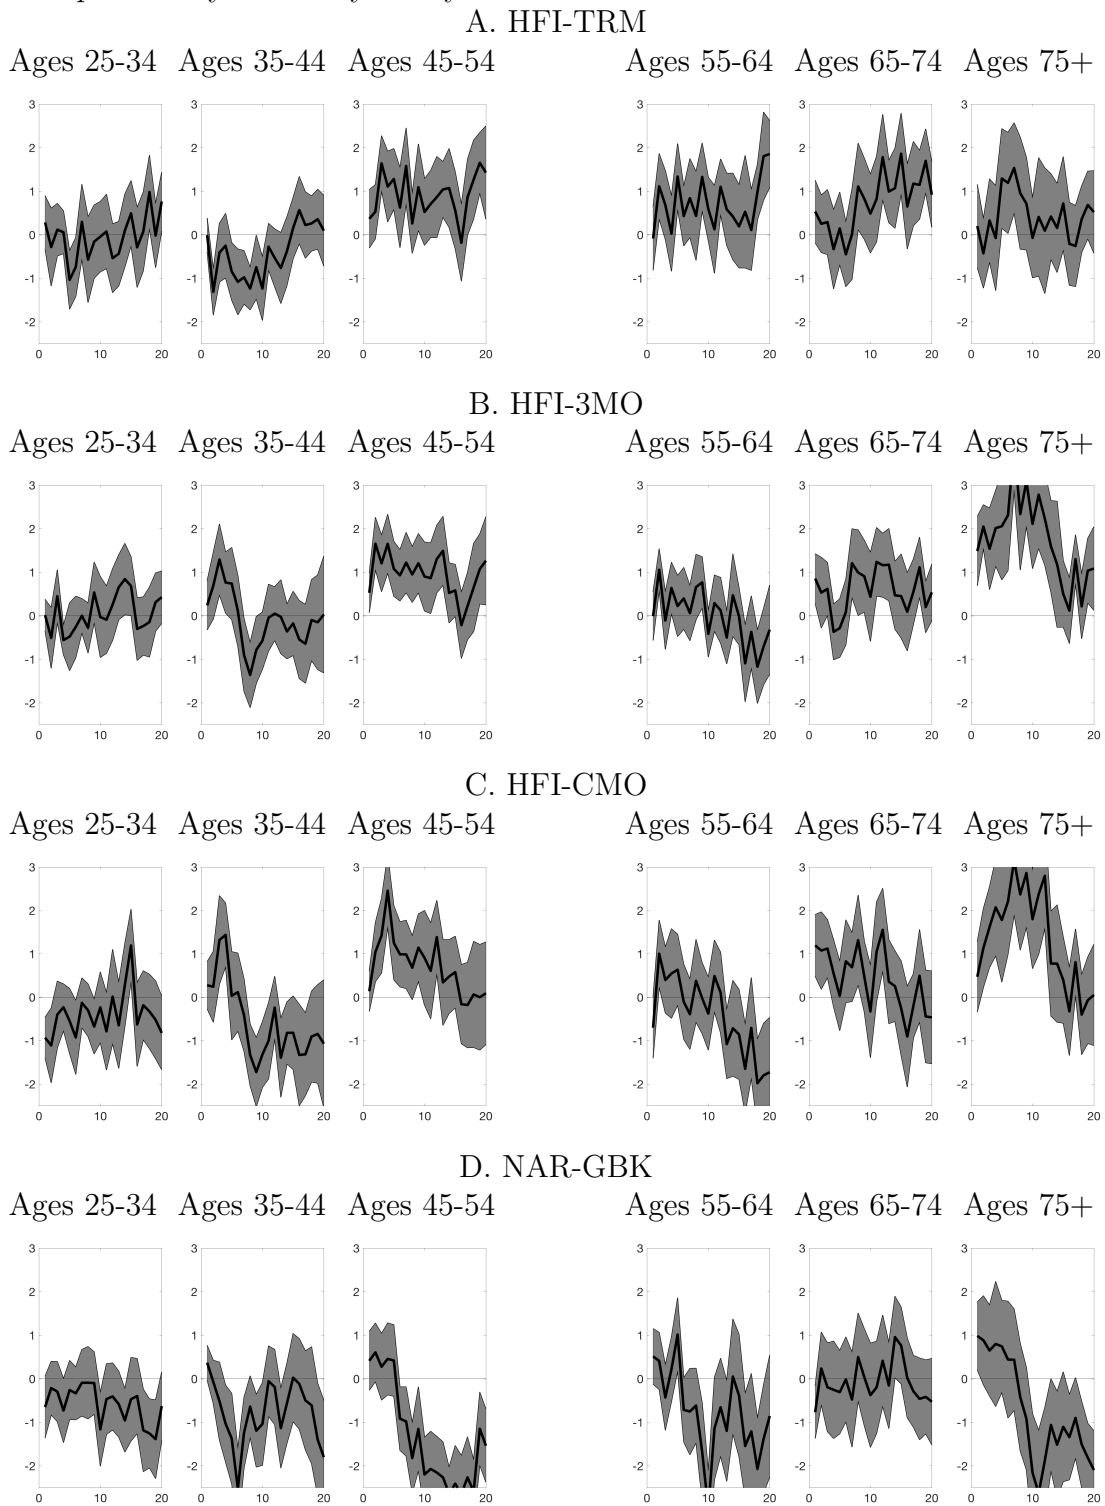
Figure B.1: Structural VAR – Cumulated Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



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  $\pm$  one standard error asymptotic confidence bands. The horizontal axis indicates quarters following the shock. The vertical axis is measured in percent.



Figure B.2: Local Projections – Cumulated Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock

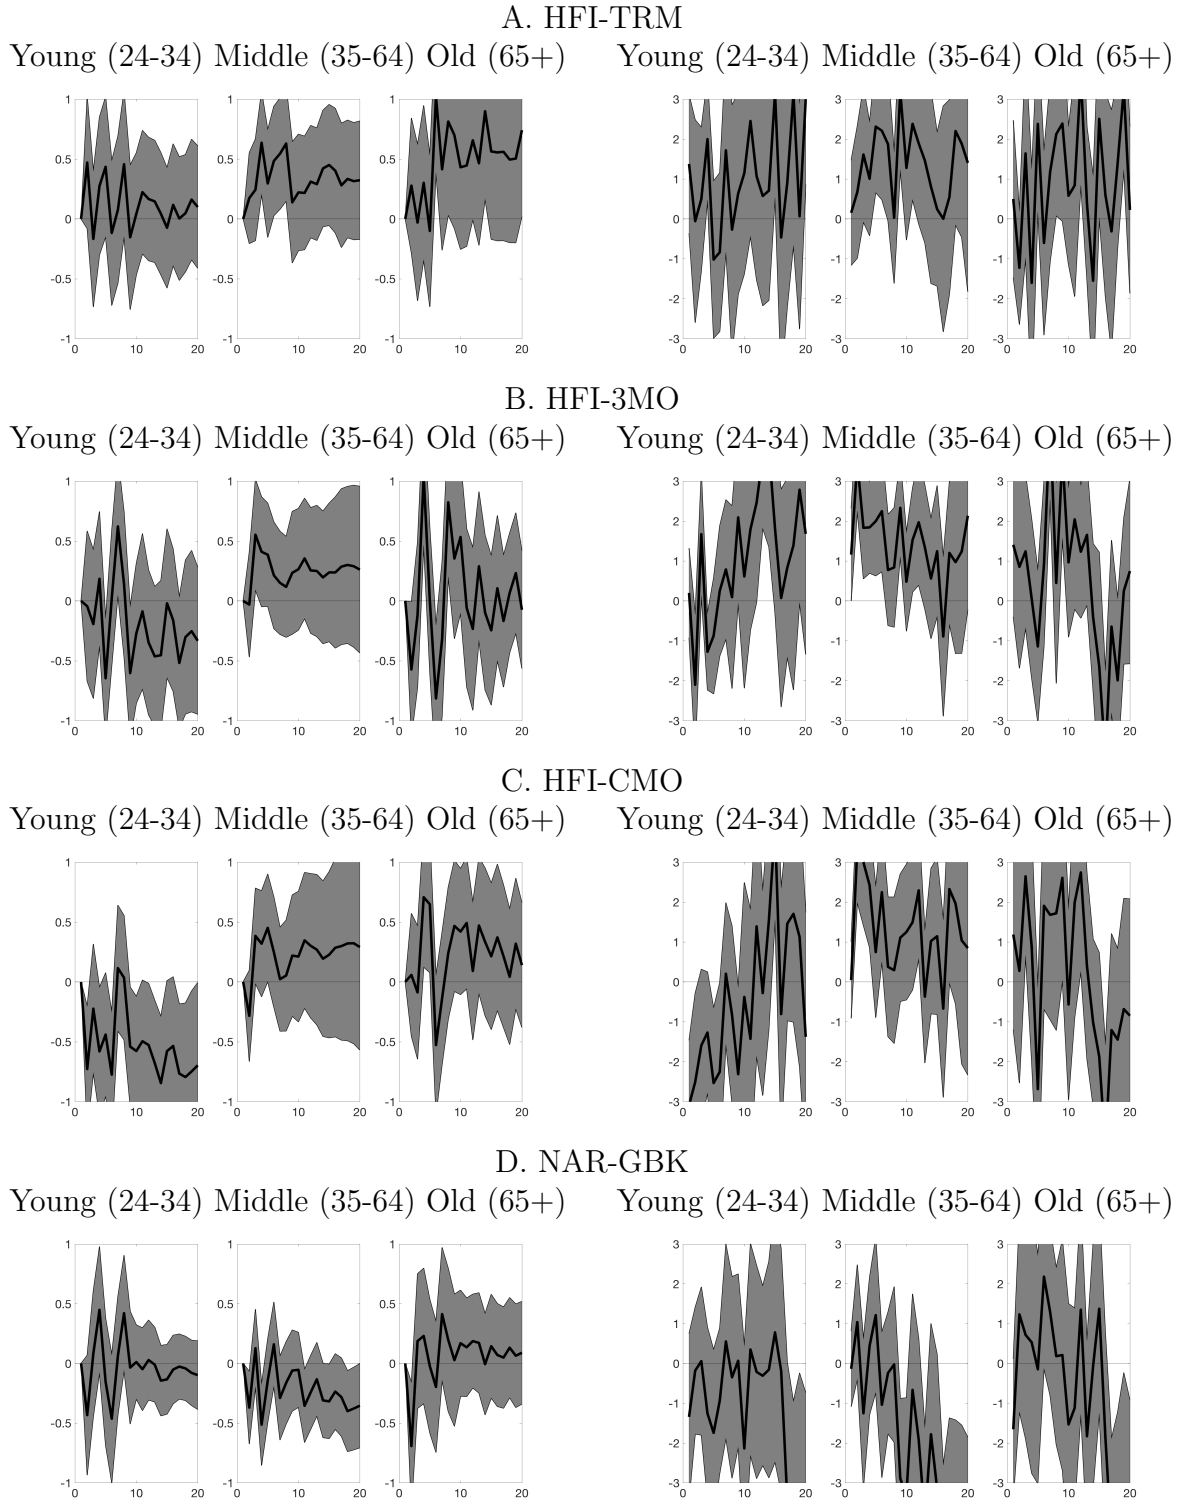


Notes: This figure plots  $-\beta_h$  from Equation (4) in the main text. Shaded areas are  $\pm$  one standard error Newey and West (1987) confidence bands. The horizontal axis indicates the number of quarters for up to five years after the shock. The vertical axis is measured in percent.

## **B.2 Durable Consumption**

For Figure B.3 we only report responses of consumption on durable goods from our structural VAR and local projections. 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.

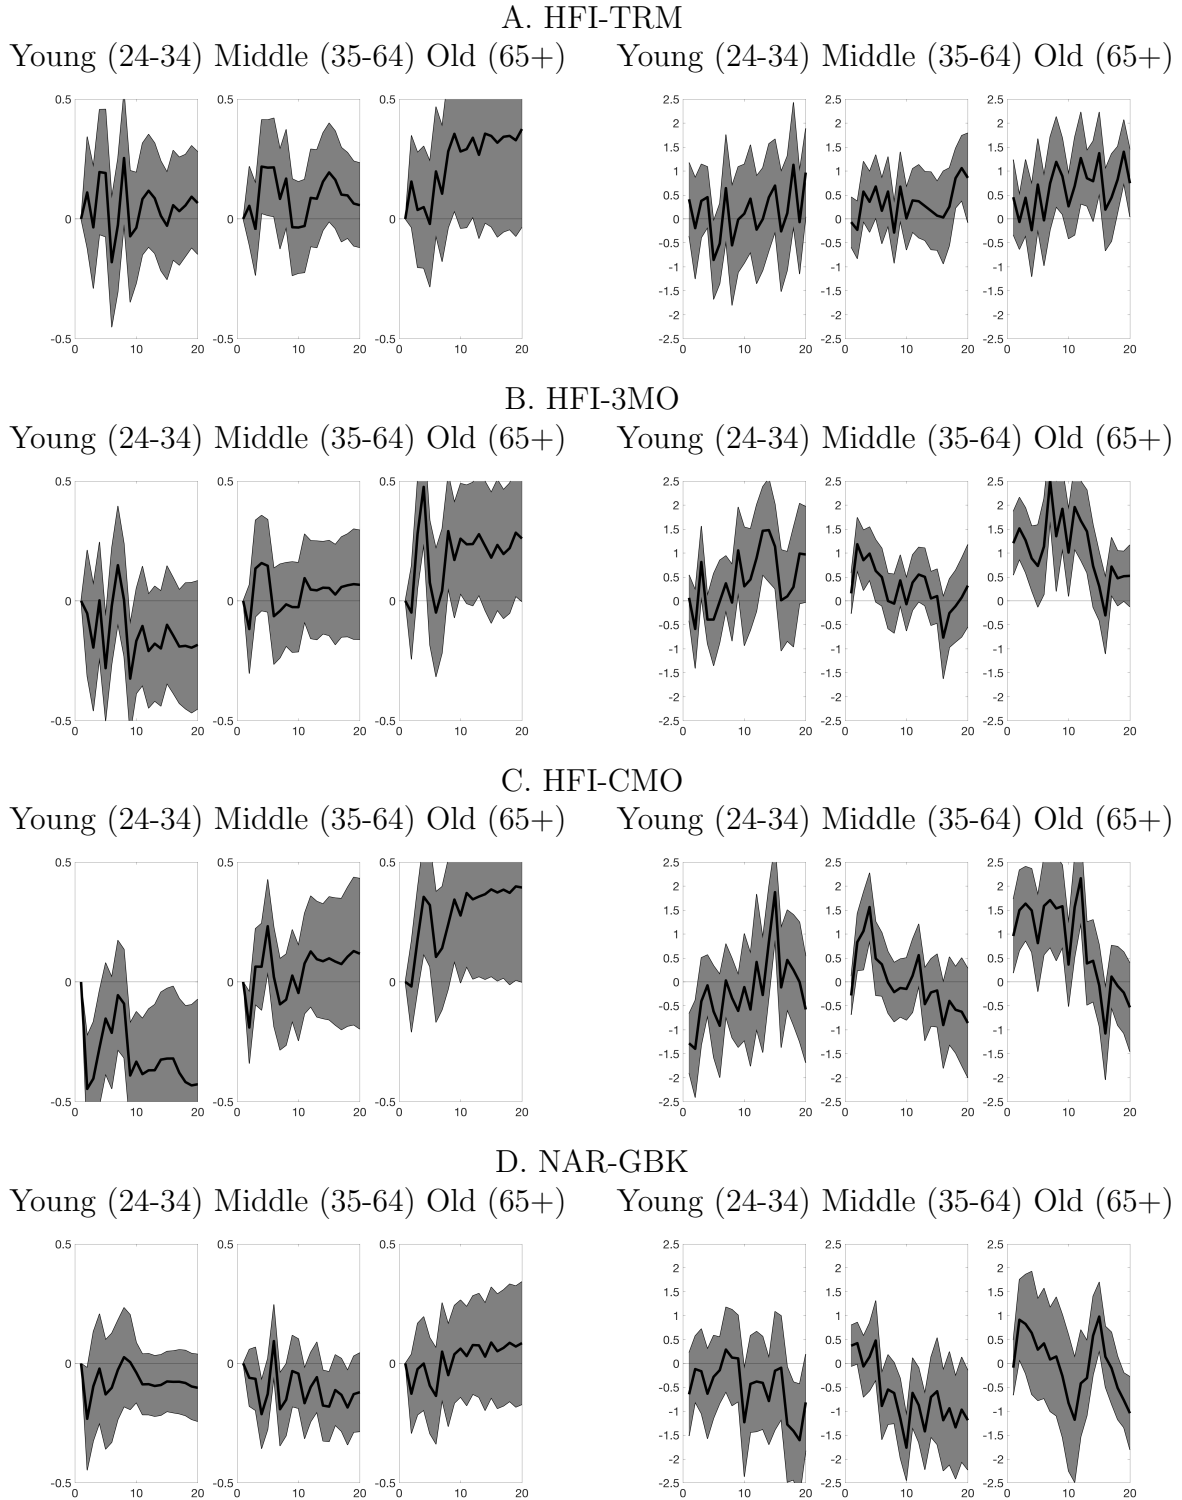
Figure B.3: Structural VAR and Local Projections – Cumulated Durable Consumption Growth  
 Impulse Response by Age Group to Expansionary Monetary Policy Shock



### **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 structural VAR and local projections in Figure B.4. These results are closely in line with our baseline results, suggesting housing expenditures are not solely driving our results.

Figure B.4: Structural VAR and Local Projections – Cumulated Consumption Less Housing Expenditures Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock



## **B.4 Alternative within Age Group Aggregation of Consumption Expenditures**

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.5 shows the responses to consumption by this measure from our structural VAR and local projections. The results are close to our baseline specification.


Figure B.5: Structural VAR and Local Projections – Cumulated Consumption (Using CEX Weights) Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock

Structural VAR

Local Projections


A. HFI-TRM

Young (24-34) Middle (35-64) Old (65+)      Young (24-34) Middle (35-64) Old (65+)




B. HFI-3MO

Young (24-34) Middle (35-64) Old (65+)      Young (24-34) Middle (35-64) Old (65+)




C. HFI-CMO

Young (24-34) Middle (35-64) Old (65+)      Young (24-34) Middle (35-64) Old (65+)



D. NAR-GBK

Young (24-34) Middle (35-64) Old (65+)      Young (24-34) Middle (35-64) Old (65+)

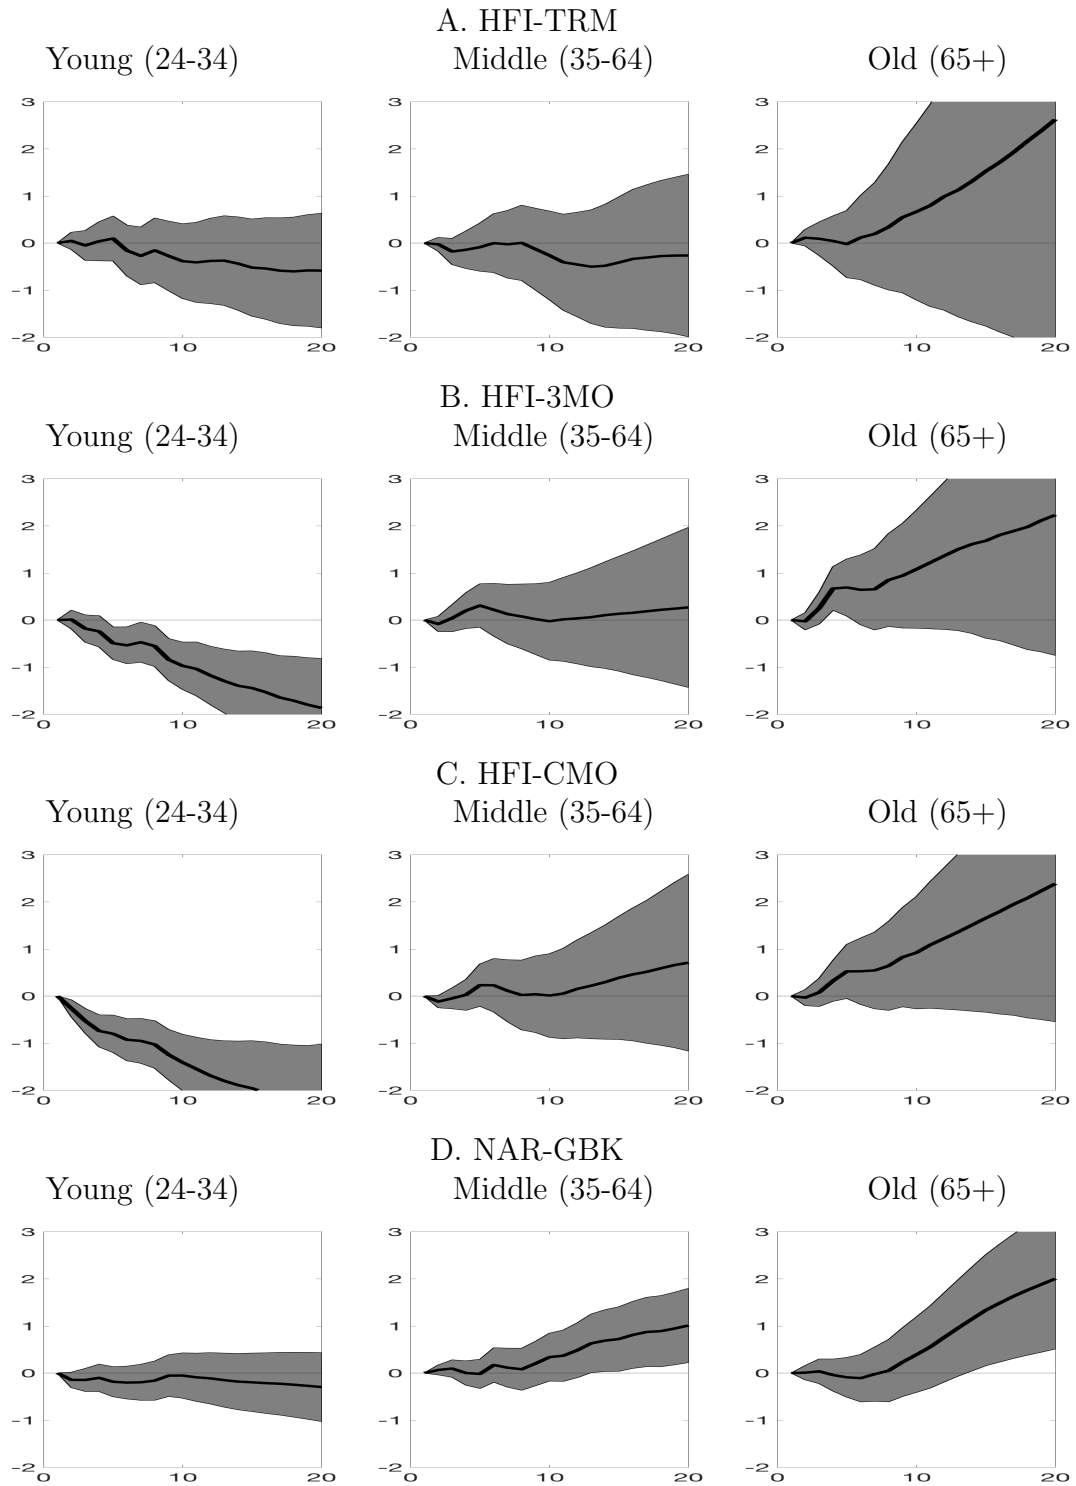


## B.5 VAR in Levels

Here, we estimate our VAR in log levels of consumption and the results are plotted in Figure B.6. The two main features of this are consistent with our baseline structural 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 NAR-GBK monetary policy shock.



Figure B.6: Structural VAR – Log Level Consumption Impulse Response by Age Group to Expansionary Monetary Policy Shock

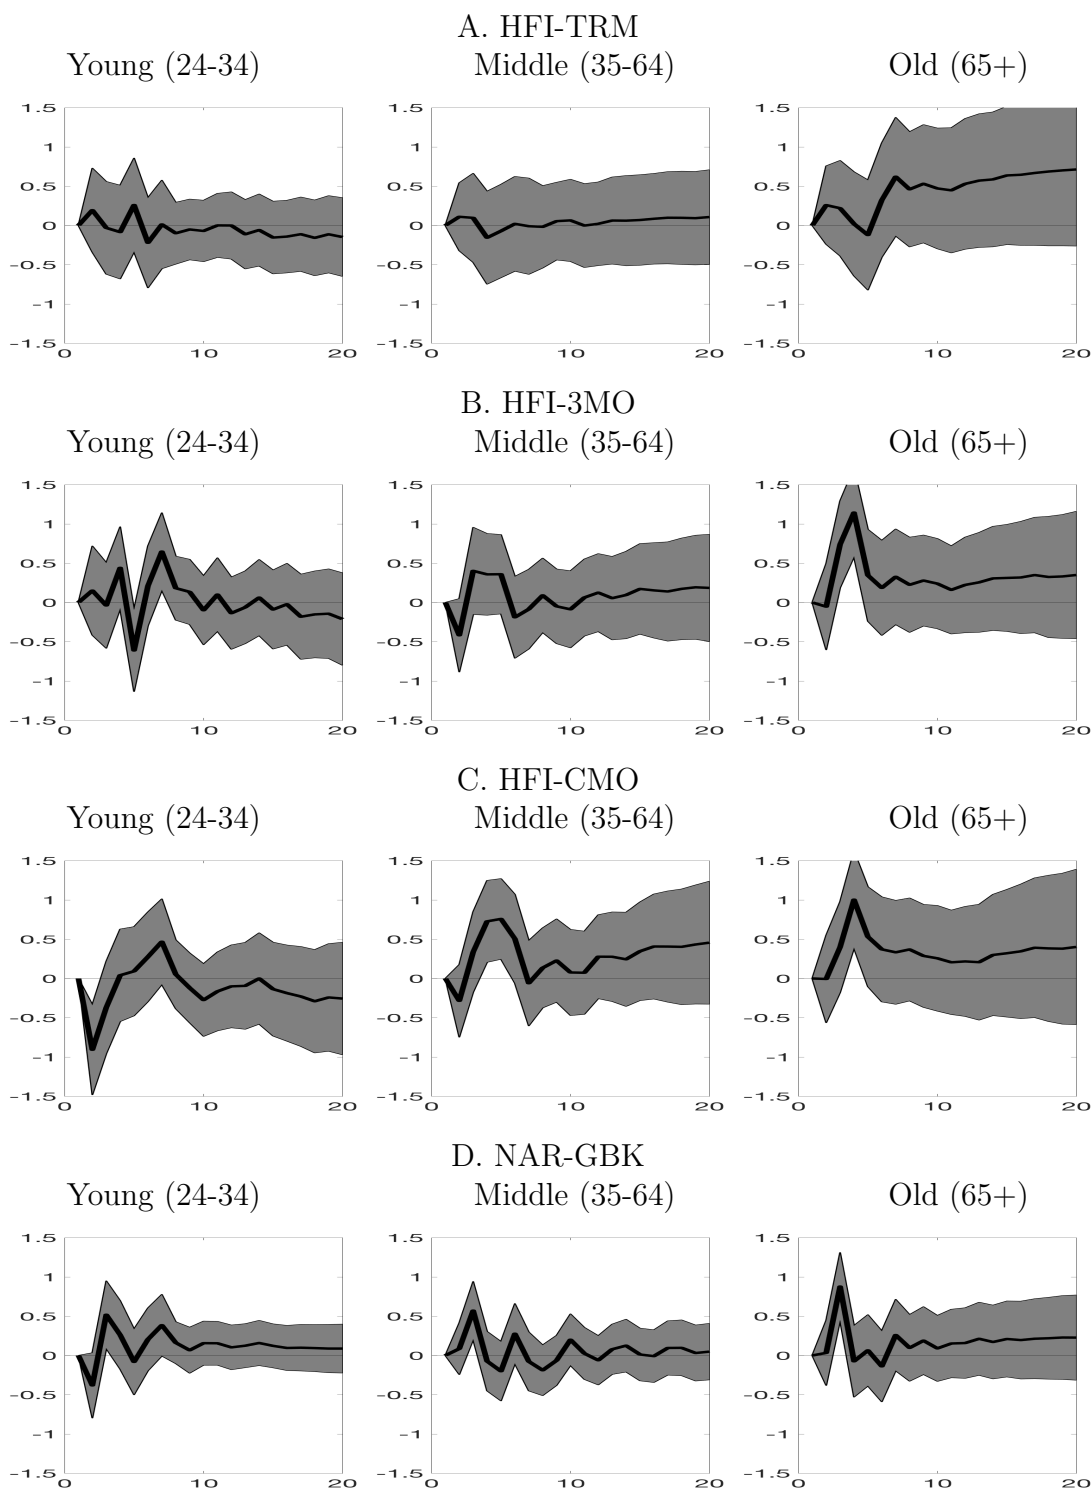


Notes: The shock is a one standard deviation decrease in the monetary policy shock series. The monetary policy shock series are normalized such that one standard deviation changes in the series match a one standard deviation change in the real federal funds rate. Shaded areas are  $\pm$  one standard error asymptotic confidence bands. The horizontal axis indicates quarters following the shock. The vertical axis is measured in percent.

## **B.6 6 Lags in the VAR**

In Figure B.7 we reduce the number of parameter estimates in our structural 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.

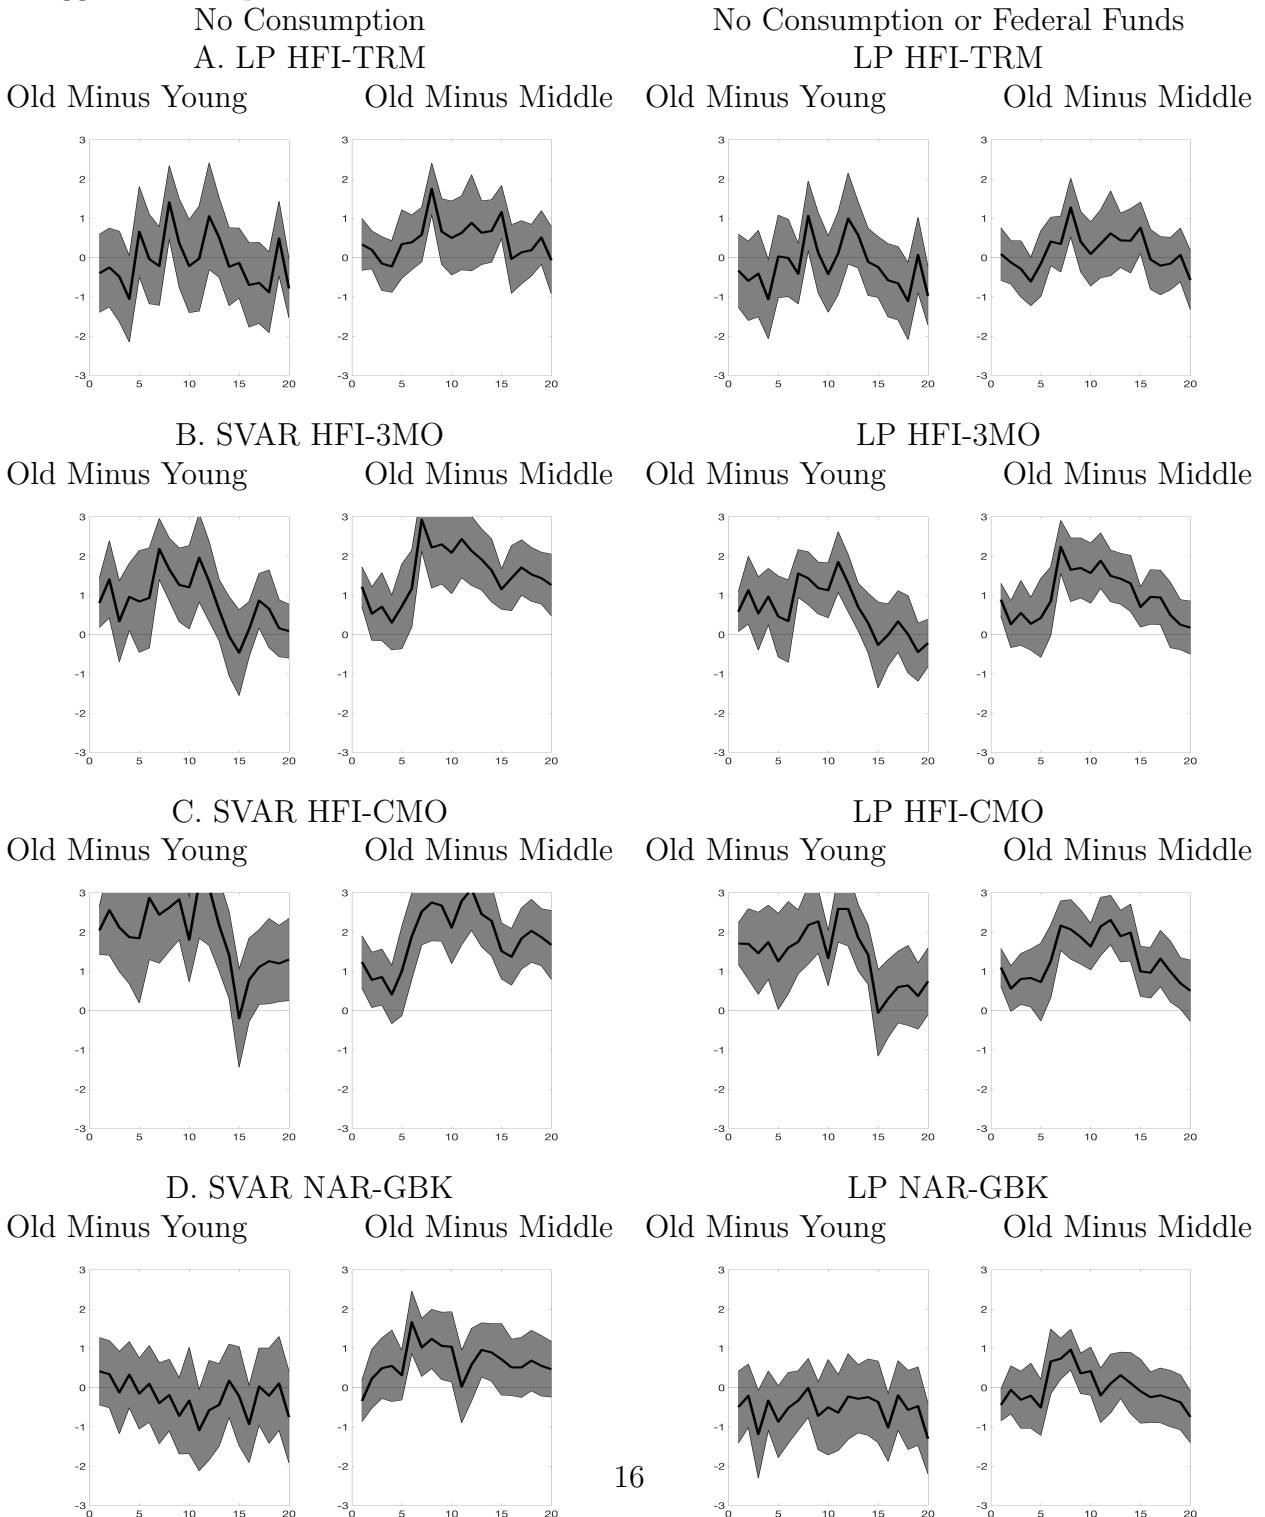
Figure B.7: Structural VAR – Cumulated Consumption Growth Impulse Response by Age Group to Expansionary Monetary Policy Shock Using 6 Lags



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  $\pm$  one standard error asymptotic confidence bands. The horizontal axis indicates quarters following the shock. The vertical axis is measured in percent.

## B.7 Local Projections Without Controls for Lagged Consumption or the Real Federal Funds Rate

Figure B.8: Local Projections (LP) – Cumulated Consumption Growth Impulse Response for Old Minus Young and Old Minus Middle to Expansionary Monetary Policy Shock without Controls for Lagged Consumption and Federal Funds Rate



## C Household (Micro)-Level Regressions

This section of the Appendix provides additional details for the estimated impulse responses by age group to monetary policy shocks based on household-level regressions, which follow the empirical strategy employed by Coibion et al. (2017) and Wong (2018). This alternative methodology generates responses qualitatively similar to the aggregate approaches (structural VAR and local projections) studied in the main paper. The micro-level analysis also finds the old to have the largest estimated response. The young generally have the second highest response, but not always.

In the CEX, households are interviewed quarterly for up to 5 consecutive quarters. There are at most 4 quarters of consumption observations per household because the first interview is only to obtain household-level characteristics but not information on consumption. Due to staggering of interview months, not all observations are based on consumption from all 3 months in the quarter. We keep only those consumption observations that have data for all 3 months in the quarter. Most households have 3 or 4 quarters of consumption observations.

The variable of interest is the percent change in quarterly consumption per household,  $i$ , at time  $t$ ,  $\Delta \ln(c_{i,t}) = \ln(c_{i,t}) - \ln(c_{i,t-1})$ . Real consumption per household is obtained by deflating the household's consumption in each specific category with the category-specific CPI and aggregating these together, as described in the main text and Appendix A. By working with growth rates, we lose one observation to differencing, leaving either 2 or 3 usable observations per household.

We further restrict the sample to households reporting non-zero expenditures on food. We also trim the sample by those consumption growth observations below -2.5 (change in log consumption) and above 2.5. If these criteria are not satisfied for even one household observation, the household is dropped from the sample. We consider the same four monetary policy shock series (HFI-TRM, HFI-3MO, HFI-CMO, and NAR-GBK). The regressions include 8 lags of the shocks to conform with the structural VAR analysis in the text.

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

$$\tilde{v}_{z,t} = \left\{ (1 - \beta) (\tilde{c}_{z,t}^\nu (1 - L_{z,t})^{1-\nu})^{1-\rho} + \beta (M_g G_{z,t} e^{n_{z,t}})^{\frac{\nu(1-\rho)}{(1-\gamma)}} [E_t \tilde{v}_{t+1}^{1-\gamma}]^{\frac{(1-\rho)}{(1-\gamma)}} \right\}^{\frac{1}{1-\rho}}. \quad (17)$$

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,

$$\tilde{a}_{z+1,t+1} M_g G_{z,t} e^{n_{z,t}} = \tilde{a}_{z,t} + r_t (\tilde{a}_{z,t} + \tilde{w}_{z,t} L_{z,t} - \tilde{c}_{z,t}). \quad (18)$$

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

$$\tilde{v}_{z,t} = \left\{ (1 - \beta) (\tilde{c}_t^\nu)^{1-\rho} + \beta (M_g G_{z,t} e^{n_{z,t}})^{\frac{\nu(1-\rho)}{(1-\gamma)}} \left[ \delta_t E_t (\tilde{v}_{t+1}^{1-\gamma}) + \frac{(1 - \delta_t) e^{n_{z,t}}}{1 - \gamma} E_t \left( \frac{\tilde{a}_{t+1}}{r_t b} \right)^{\nu(1-\gamma)} \right]^{\frac{(1-\rho)}{(1-\gamma)}} \right\}^{\frac{1}{1-\rho}}, \quad (19)$$

and

$$\tilde{a}_{z,t+1} e^{n_{z,t}} = \tilde{a}_{z,t} + r_t (\tilde{a}_{z,t} + \tilde{w}_{z,t} - \tilde{c}_{z,t}). \quad (20)$$

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

$$A_{z,t+1} = (1 + r_t) (A_{z,t} + L_{z,t} W_{z,t} - C_{z,t}). \quad (21)$$

## E Variations on the Model

In this section, we report results from alternative specifications of the model.

Panel A of Figure E shows consumption responses when the asset is a short-term (one-period) bond. Other model features are the same as those in the text,  $\gamma = 12$ ,  $\psi = \rho^{-1} = 1.4$ ,  $\nu = 0.5$ , and the bequest motive is present. Size-ordering of the consumption responses goes young > old > middle, which is inconsistent with our empirical estimates.

Panel B shows consumption responses when labor supply is fixed, again with  $\gamma = 12$ ,  $\psi = \rho^{-1} = 1.4$  and the bequest motive. When households make only the consumption-saving decision, size-ordering of consumption responses is young > old  $\simeq$  middle, which again is inconsistent with our empirical estimates.

Panel C shows consumption responses under constant relative risk aversion utility with the labor decision and bequest motive,  $\gamma = 12$ ,  $\nu = 0.5$ . Here, consumption from all age groups respond in the wrong direction.

Figure E.1: Variations on the Model





Figure E.2 shows relative consumption responses when the intertemporal elasticity of substitution is  $\psi = \rho^{-1} = 0.2$  instead of 1.4.

Figure E.2: Alternative Preference Parameter Settings  
 $\gamma = 12, \psi = \rho^{-1} = 0.2, \nu = 0.5$



While the age-group consumption response ordering is consistent with our estimates, a further breakdown across high and low wealth young households does not conform to the data. As seen in Figure E.3, consumption of the wealthy young increases whereas consumption of the poor young declines initially. Similarly, labor response of the poor young lies above that of the wealthy young. An intertemporal elasticity of substitution greater than 1 is needed to get the poor young consumption response to lie above the wealthy young response.

Figure E.3: Alternative Preference Parameter Settings  
 $\gamma = 12, \psi = \rho^{-1} = 0.2, \nu = 0.5$

