NBER WORKING PAPER SERIES

DEMOGRAPHIC EFFECTS ON THE IMPACT OF MONETARY POLICY

John V. Leahy
Aditi Thapar

Working Paper 26324
http://www.nber.org/papers/w26324

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

We would like to thank V.V. Chari, Jordi Gali, Simon Gilchrist, Christopher House, Pablo Ottonello, Lasse Pedersen, Morton Ravn, Nick Sly, Chris Waller, Jonathan Willis, and Arlene Wong for helpful comments and discussions, and Jessica Meadows and Amanda West for valuable research assistance. Portions of this paper were completed while Dr. Thapar was visiting the Federal Reserve Bank of Kansas City as part of CSWEP's Summer Economic Fellows Program. We are grateful to the National Science Foundation for their support under grant SES-1919362. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Demographic Effects on the Impact of Monetary Policy
John V. Leahy and Aditi Thapar
NBER Working Paper No. 26324
September 2019
JEL No. E32,E52,J11

ABSTRACT

We study whether the effects of monetary policy are dependent on the demographic structure of the population. We exploit cross-sectional variation in the response of US states to an identified monetary policy shock. We find that there are three distinct age groups. In response to an increase in interest rates, the responses of private employment and personal income are weaker the greater the share of population under 35 years of age, are stronger the greater the share between 40 and 65 years of age, and are relatively unaffected by the share older than 65 years. We find that all age groups become more responsive to monetary policy shocks when the proportion of middle aged increases. We provide evidence consistent with middle aged entrepreneurs starting and expanding businesses in response to an expansionary monetary shock.

John V. Leahy
Department of Economics and
Gerald R. Ford School of Public Policy
University of Michigan
3308 Weill Hall
735 S. State St. #3308
Ann Arbor, MI 48109
and NBER
jvleahy@umich.edu

Aditi Thapar
Dept. of Economics
Univ. of Michigan
611 Tappan Ave
Ann Arbor, MI 48109
athapar@umich.edu

1 Introduction

The world is undergoing a period of demographic change. Populations are aging in much of the developed world, especially Japan and Continental Europe. On the other extreme, over 50% of the population of India is below the age of 25.[1] In China, the on-again-off-again one child policy is likely to bring large fluctuations in the age distribution. Given the life-cycle patterns in many economic activities such as saving, home purchases, education, and retirement, it is of interest to understand how changes in the age distribution affect the performance of the tools of economic policy. In this paper we focus on the relationship between the age structure of the economy and monetary policy.

To understand how demographics affect monetary policy, we need a sample of economies with differing demographic structures and we need a measure of monetary policy. We use the US states as our laboratory. The states differ in their age structure, and data is available on a consistent basis for a wide range of variables. Since the states share a single monetary policy, we do not have to control for policy differences across states or worry that policy might be responding to the demographic structure of a given state. This allows us to focus on how changes in the age structure influence the response to this common policy. Of course, states differ in other ways, some of which may be correlated with age structure, and we will attempt to control for these differences in what follows. Heterogeneity, however, is present, more or less, in nations as well.

For monetary policy, we follow the high-frequency approach to the identification of monetary policy shocks and consider changes in the Federal Funds rate over a short window around policy announcements.[2] We ask how the age distribution in a state affects the response of state income and employment to a monetary policy shock. We find that monetary policy is more effective – in the sense that income and employment respond more to a given change in the interest rate – the greater the share of the middle-aged population, defined as those aged 40 to 65, and less effective the greater the share of the young, defined as those aged 20 to 40. The share of the population over 65 has no discernible effect on the effectiveness of monetary policy.

[1] 2011 Census, http://www.censusindia.gov.in/2011-Common/CensusData2011.html

[2] See Gürkaynak, et al. (2005), Gertler and Karadi (2015), and Nakamura and Steinsson (2018) for prominent examples of this approach.

These effects of demography appear sizable, and comparable in size to the effects of a monetary policy shock on aggregate income. In the baseline analysis, a one standard deviation increase in the fraction of 30-34 year olds reduces the impact of a monetary policy shock by approximately 40%, whereas an increase in the fraction of 50-54 year olds increases the impact by a similar magnitude.

The basic pattern in which the middle aged amplify and the young dampen the effects of monetary policy appears robust. It holds for a variety of definitions of the monetary policy shocks and for a variety of empirical specifications. We investigate altering the start and end dates of the sample, the removal of small states, and aggregation to census regions.

Having established the robustness of our results we turn our attention to potential mechanisms. We begin by asking whether our results reflect demographics or some other state characteristic correlated with demographics. We investigate two possibilities in detail: state industrial composition and state income. If a large middle age population were correlated with a large manufacturing sector, the greater response of manufacturing to monetary policy shocks could explain the greater sensitivity to interest rates. Alternatively, Cravino, Lan and Levchenko (2018) have argued that richer individuals consume a greater share of sticky price goods. If a large middle age population were correlated with high income, stickier prices could then explain the greater response of middle aged states to monetary shocks. We find no such correlations and therefore no evidence that either industrial composition or state income is driving our results. We also consider a number of other potential correlates including demographic variables such as gender, race, and education and housing variables such as home ownership and house prices. Again the basic pattern survives.

We then investigate whether the response of employment to a monetary policy shock is due to hiring within the state or migration between states. If the response were due mainly to migration, then this would suggest caution in applying the results to sovereign nations as migration between states is likely easier than migration between nations. We find that the response of unemployment mirrors the response of employment, suggesting an important role for internal adjustment. We look at the response of state population and find some evidence that the age structure of the state affects the response of migration to a monetary policy shock. The response of migration to the age structure, however, appears to be largely orthogonal to the responses of income and employment.

Outflows in response to tight policy are larger in states with more 35-40 year olds and inflows are larger in states with more elderly. Both are age groups that do not greatly affect the response of income or employment.

A greater proportion of middle aged residents could amplify the effects of monetary policy in a number of ways. On one hand, the income and employment of middle aged residents could respond more to monetary policy shocks. Having more middle aged residents would then explain the greater aggregate response. On the other hand, having more middle aged residents could increase the responsiveness of the income and employment of all age groups. We find evidence for the second channel. We look at the response of income and employment broken down by age. We find that the incomes and employment of all age groups follow the same basic pattern: dampened if there are more young in the state and amplified if there are more middle aged. It is not so much being middle aged that matters as being surrounded by middle aged people.

Finally we suggest a mechanism by which a greater number of middle aged residents amplifies the effects of monetary policy on all age groups. The mechanism is motivated by the fact that the middle aged are more likely to be entrepreneurs. Selgado (2018) documents that entrepreneurs tend to be older than workers and that the share of entrepreneurs in the 35-60 age group is larger than the share of entrepreneurs in the 22-25 age group. Azoulay, Jones, Kim and Miranda (2019) argue that successful entrepreneurs are middle-aged not young. If a larger share of middle aged residents implies a larger number of potential entrepreneurs, then the elasticity of business formation with respect to a decline in interest rates is likely larger in states with a greater share of the middle aged population. We find that business formation does in fact respond more strongly the greater the share of the middle aged. We find evidence for a multiplier, by which the initial expansion increases business income furthering the incentive to invest. We also find evidence of financial frictions. A greater fraction of middle aged increases the response of small business employment to a monetary policy shock. The hypothesis is then that a monetary expansion increases the returns to starting a small business both by lowering the cost of borrowing and by increasing the profitability of a business once it is started. The greater the proportion of the middle aged the greater the response to these incentives. All income groups then benefit from this expansion.

In Section 2 we describe recent background literature that is relevant to our investigation. In Section 3 we introduce our empirical strategy. We present our main results in Section 4. We discuss interpretations of our results in Sections 5 and 6. Section 7 concludes.

2 Related Literature

This paper contributes to the small, but growing, literature on the relationship between monetary policy and the demographic structure of the economy.[3] Several recent papers study how monetary policy shocks impact the consumption of different age groups. Sterk and Tenreyro (2018) find that monetary policy shocks have greater effects on the durable consumption of the young than they do on the middle aged and old. Wong (2018) also finds that the consumption of the young is more sensitive to monetary policy shocks and argues that young homeowners, who either obtain a new loan or refinance, are the primary drivers of this response. Cloyne, Ferreira and Surico (2018) argue that consumption of households with mortgages and those who are renters rise after expansionary shocks, while outright owners don't respond. They argue that the first group (households with mortgages and renters) are typically younger, while the second group (outright homeowners) are typically older. These studies all focus on household responses. Using aggregated data, Berg, et al. (2019) argue instead that older households as a group have a higher consumption response than younger households.

We find something very different, namely that monetary policy is more effective the greater the number of middle aged and less effective the greater the number of young. One potential reason for this difference is that we look at a different outcome variable. In their household level analysis, Sterk and Tenreyro, Wong, and Cloyne, Ferreira and Surico, all look at the effects of monetary policy on consumption. In our state level analysis, we look at the effects on income and employment since quarterly US state consumption data is only available after 2005. It is very likely that consumption and income respond in different ways to monetary policy. The relationship between an individual's consumption and income depends to a large extent on the individual's position in the life cycle.

[3] There is a related literature considering the distributional effects of monetary policy. Prominent examples include Doepke and Schneider (2006), Coibion, Gorodnichenko, Kueng and Silvia (2017), Kaplan, Moll, and Violante (2018), Cravino, Lan and Levchenko (2018), and Auclert (2019). Most of this literature focuses on how monetary policy affects income inequality.

Liquidity constraints are more likely to affect the young, and the middle aged are generally in a better position to smooth consumption (Gourinchas and Parker (2002)). It is therefore likely that the consumption of the young responds more to monetary policy shocks, while income and employment depend on the proportion of the middle aged.

Another potential reason that our results differ from the previous literature is that our question is somewhat different. Rather than ask how monetary policy affects different age groups, we ask how changes in the age-composition of the economy alter the impact of monetary policy shocks on aggregate economic performance. If aggregate responses are merely the sum of individual responses, then the answers to these two questions would be the same. One of our principle findings, however, is that the effects of age composition appear to work through general equilibrium effects, and that these general equilibrium affects hit all age groups in a similar manner. What matters for the response of the income of the young in Idaho is the overall age distribution in Idaho and not the fact that the group is young. For both of these reasons we see our findings as consistent with and complementary to the existing literature.

One paper that considers how the age-composition of the economy alters the impact of monetary policy is Imam (2015). Imam assesses the impact of monetary policy on unemployment and inflation as economies age. He first estimates a time-varying coefficient Bayesian VAR with stochastic volatility for a panel of the large OECD economies, Canada, Germany, Japan, U.K., and U.S. Monetary policy shocks are identified by a Cholesky decomposition with money ordered last. Using the estimated impact of monetary policy on unemployment and inflation, he then proceeds to estimate the impact of aging on the variables of interest. He uses the old-age dependency ratio, the ratio of the number of people older than 65 years to those aged 15 to 64 years, as his demographic variable. He therefore has only two age groups. He finds that policy has become less potent as the fraction older than 65 has risen. He does not consider how the middle aged might differ from the young.

Several papers, including Fujiwara, Ippei, and Yuki Teranishi (2006), Kantur (2013), Carvahlo, Fererro, and Nechio (2016), Kara and von Thadden (2016), and Sterk and Tenreyro (2018) calibrate dynamic general equilibrium models to study the effects of demographic change on monetary policy. Most of these papers build on the model of Gertler (1999) in which individuals progress through two life stages: the young work and the old are retired. Our results suggest the need for a richer demographic structure. At a minimum one needs to distinguish between young and middle aged workers, in addition to the retired.[4]

Our paper contributes to a literature that uses a panel of US states to answer aggregate questions. Most of these papers study fiscal policy. Shoag (2010), Nakamura and Steinsson (2014), and Chodorow-Reich (2018) analyze fiscal policy multipliers at the state level. In concurrent work, Basso and Rachedi (2018) study the effects of fiscal shocks interacted with the proportion of the state population in three age bins, and find that the fiscal multiplier is larger the greater the number of young people in the state.

We are not the first to argue that the share of middle aged might have important consequences for economic activity. Freyer (2007) finds that an increase in the proportion of workers between 40 and 49 increases productivity growth in a cross section of OECD countries. Azoulay, et al. (2019) show that founders of successful businesses tend to be middle-aged. Selgado (2018) argues that entrepreneurs tend to be middle aged.

3 Empirical Implementation

We treat the US states as a panel of economies. There are several advantages to using the states as a laboratory. The primary advantage is that the US is a monetary union so that all states are subject to the same monetary policy. This means that we can remove the first-order effect of monetary policy using time fixed-effects and focus on the interaction of monetary policy shocks with local demography. A common monetary policy also removes many sources of heterogeneity. We do not have to worry about heterogeneity across states in monetary policy rules. We do not have to worry about monetary policy responding to the demographic makeup of each state. We do not have to worry about differences that might arise from differences in the identification of monetary policy shocks across different nations. The states also have a homogeneous legal and institutional

[4]A couple of recent papers, Gagnon, Johanssen and Lopez-Salido (2016) and Eggertsson, Mehrotra and Robbins (2017), use a richer demographic structures to understand the decline in output growth and the real interest rate. These papers do not study the economy's response to monetary policy.

structure. The problems that remain, such as immigration, trade, and heterogeneity of industrial structure are present, more or less, in nations as well, and we will attempt to control for some of these influences in what follows.

In the remainder of this section we discuss the empirical implementation of our approach. We begin by describing the empirical specification, followed by a discussion of how we identify monetary policy shocks, and a description of the data used in this paper.

3.1 Specification

We use variation across U.S. states and across time in the share of population in different age groups to estimate how the population structure affects the response of income and employment to monetary policy shocks. We consider the following specification:

$$\Delta log[X_{s,t+i}] = \alpha^{i,a}\epsilon_t^m p_{s,t}^a + \phi^a p_{s,t}^a + \sum_{j=0}^{J} \beta^j Z_{s,t-j} + \gamma_s + \delta_t + u_{s,t}. \quad (1)$$

Here s indexes the U.S. states, a indexes the age-bins, and t indexes time. In our baseline specification we consider all 50 US states plus the District of Columbia. The age bins run every five years from age 20 to age 69 and then we group the population over 70 into an additional bin. By considering five year age bins we do not place any a priori restrictions on what age groups matter for monetary policy and how these age groups matter. The unit of time is a quarter.

The dependent variable is the log change in X_s between period $t-1$ and period $t+i$. Our main results take X to be either nominal personal income or private employment. ϵ_t^m is an exogenous shock to U.S. monetary policy. $p_{s,t}^a$ is the percentage of the population of state s that belongs to age-group a at time t. $Z_{s,t-j}$ is a vector of additional control variables. We include the four quarter growth rate in the population in the state in most specifications to control for trends in population growth. The state effects γ_s capture persistent differences in age or industrial structure across states, and time fixed-effects δ_t account for factors that are common to all states such as interest rates, the exchange rate and aggregate inflation. $u_{s,t}$ is the error term. Note that since we include time effects, we do not include the money shock ϵ_t^m separately in the regression.

All percentages such as income growth, the monetary shock, and the population shares are expressed in basis points.

We run this panel regression separately for each horizon i and each age-bin a. The coefficients of interest are the $\alpha^{i,a}$ which capture how the effect of a shock to monetary policy in period t on X_{t+i} varies with the proportion of the population in age-group a. Since our monetary policy shocks represent shocks to interest rates, positive shocks reflect contractionary monetary policy. Positive values of $\alpha^{i,a}$ therefore imply that an increase in the proportion of the population in age-group a reduces the effect of a shock to monetary policy on X. Negative values of $\alpha^{i,a}$, on the other hand, imply that as the population in age-group a increases the contractionary effects of an increase in interest rates become stronger.

For any given age-group $\bar{a}$, estimating equation (1) for each $i = 1, 2, ...H$ indicates how an increase in the proportion of the population in group $\bar{a}$ affects the response of X to a monetary policy shock over time. Alternatively, for any given horizon $\bar{\imath}$, estimating equation (1) indicates how an increase in the proportion of the population in group a affects the response of monetary policy on X at horizon $\bar{\imath}$ for different age groups.

3.2 Monetary policy shocks

In our baseline specification, we follow the high-frequency approach to the identification of monetary policy shocks employed by Gürkaynak, et al. (2005), Gertler and Karadi (2015), and Nakamura and Steinsson (2018).[5] As highlighted by Nakamura and Steinsson (2018), there are some concerns regarding the interpretation of these high-frequency shocks. We therefore consider the robustness of our results to other shocks in Section 4.2.1. We provide a detailed explanation of how we construct our shocks in Appendix B. Here we provide a brief overview.

There are two steps in the construction. First, we obtain raw shocks from the movement in the federal funds rate implied by the current month federal funds futures contract in a relatively short window of time around the FOMC announcements. We consider the change in the futures rate from 15 minutes before the announcement to 45 minutes after the announcement. By considering a narrow window around the FOMC announcement, one can be reasonably certain that no other

[5] Interest rate futures have been used to identify monetary policy shocks by Rudebusch (1998), Kuttner (2001), Soderstrom (2001), and Thapar (2008).

news caused the change in the futures rates.

Second, following Ottonello and Winberry (2018), we aggregate the high-frequency shocks to a quarterly frequency using a weighted moving average of the shocks in the current quarter and the previous quarter. These weights take into consideration the fact that employment and personal income are the result of outcomes and decisions that accumulate over the entire quarter and that the monetary policy shock can only affect the subset of these decisions that are made subsequent to the shock. The weights are chosen so that an announcement at the end of quarter $t-1$ is essentially the same as an announcement at the beginning of quarter t.[6]

The shocks are illustrated in Figure 11 in Appendix B. The identified shocks are small but not insignificant. The average high-frequency shock over the 1990-2008 sample is close to zero (-5 basis points) with a standard deviation of 11 basis points. Approximately 40% of shocks were contractionary while 60% are expansionary, and there are 70% more shocks than actual changes in the federal funds rate. Very often the absence of a change in interest rates is itself a shock.

Since our shocks only represent the unanticipated component of monetary policy at the time of the meetings, they may reflect only a fraction of exogenous monetary policy changes. This makes it difficult to judge the quantitative significance of the coefficients $\alpha^{i,a}$. To establish a benchmark we look at the effects of monetary shocks on the aggregate economy. We follow Ramey (2016) and estimate,

$$log[X_{t+i}] = \alpha^i \epsilon_t + \sum_{j=1}^{J} \beta^j \epsilon^{hf}_{t-j} + \sum_{j=1}^{J} \gamma^j ffr_{t-j} + \gamma_s + \sum_{j=1}^{J} \delta^j Z_{t-j} + u_t. \quad (2)$$

where, $i = 1, 2, ..., H$ represents the horizon, X is aggregate US real GDP, ϵ_t is our monetary policy shock, ffr represents the level of the federal funds rate, and Z (following Ramey) includes lags of X, current and lagged real GDP, the GDP deflator, and the unemployment rate. To account for serial correlation in the residuals, we estimate Newey-West standard errors. Note that all of these controls would be subsumed into the time effects in our baseline specification (1).

Figure 1 depicts the response of U.S. real GDP, real personal income, and private employment to a one percentage point shock to monetary policy over a horizon of four years. We normalized the

[6]Note that this second step introduces some serial correlation in the shocks which we will have to control for in the estimation.

Figure 1: Response to a Monetary Policy Shock

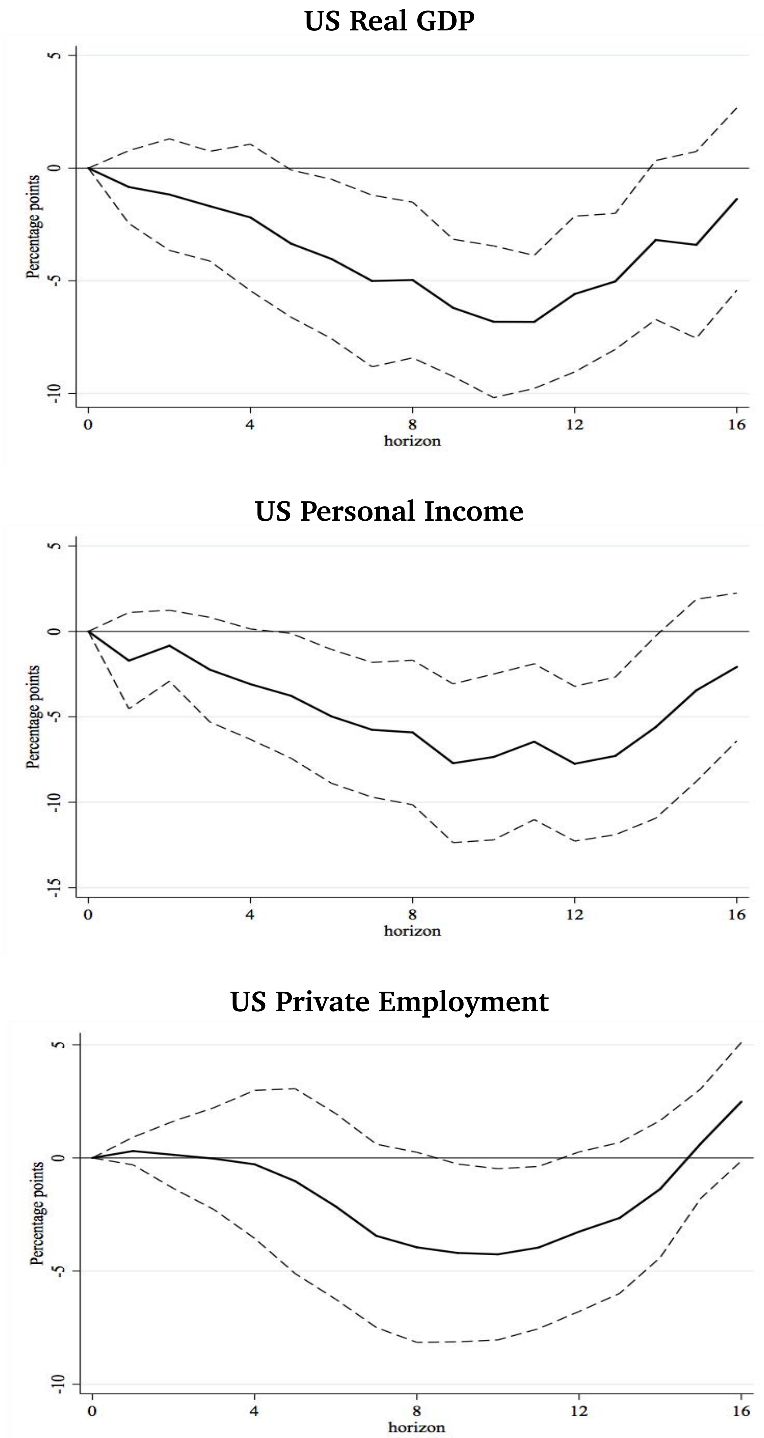

The solid line represents the response to a positive one-standard deviation increase in the federal funds rate. The dashed line represents the 95% confidence intervals given Newey-West standard errors.

shock to one percentage point because the units are easy to interpret and relate to the discussion that follows. It should be noted that this scaling greatly magnifies the effects of monetary policy shocks relative to the data. The standard deviation of shocks in our sample is an order of magnitude smaller at 0.11.

The impulse responses in Figure 1 look like the effects of a monetary policy shock. As anticipated, this shock leads all three variables to decline. The decline in real GDP and real personal income become statistically significant 5 quarters after the shocks. The reduction in GDP reaches a maximum two to three years after the shock. We will compare the $\alpha^{i,a}$ in (1) to the β^{j}in (2) as a means of judging the quantitative significance of the interaction between demographics and monetary policy.

3.3 Data

The federal funds futures market began in late 1988. Given that there was likely a learning period, we begin our sample in 1990. To avoid the zero lower bound on interest rates, we end in 2008. We consider alternative start and end dates in the robustness section.

We require state-level data on various macroeconomic variables, population estimates by state, and monetary policy shocks to the US economy. We describe each of these in turn below.

In an ideal world, we would have quarterly state-level GDP and consumption data. Unfortunately, this data is available only from 2005 on. We use the next best alternative, namely, state personal income. Total state personal income, wages and salaries, dividends interest, and rental income, and non-farm proprietors' income are obtained from the Regional Economic Accounts program at the Bureau of Economic Analysis. We use nominal personal income since there are no state-level price indices available for our sample period.[7] Deflating the nominal personal income of each state by a national price index would not alter our results as the aggregate price index would simply get soaked up by the time fixed effects in the estimation.

All state level employment data used in this paper is from the Current Employment Statistics program conducted by the Bureau of Labor Statistics. State level employment data is available

[7] Since 2008, the BEA has begun publishing annual Regional Price Parities which compare prices in states relative to overall prices for the U.S.

beginning in 1990. The data is published at a monthly frequency, which we average to a quarterly frequency. Annually, the monthly estimates are revised based on a complete count of jobs derived from tax filings of employers. The data we use is thus not affected by any sampling issues. We use data on private employment, as well as employment in three main industries: manufacturing, construction, and financial activities. For employment in the construction industry, data is available for only 44 states.[8]

State population data is obtained from the Census Bureau. Over our sample period, the data is available on an annual frequency.[9] We do not interpolate the data and assume that population remains constant within the calendar year. $p^a_{s,t}$ is defined as the number of individuals in the age-bin a in state s at date t, divided by the total population in the state at that time, times 100.

The monetary policy shocks were provided to us by Ottonello and Winberry.

4 Results

4.1 Baseline Results

Our baseline specification takes X to be either state personal income or private employment and Z to be the change in state population between t and $t-4$. The monetary policy shocks are the high frequency shocks discussed in section 3.2.

Figure 2 presents estimates of $\alpha^{i,a}$ for $i = 6$ in equation (1), when X is state personal income and state private employment, respectively. Coefficients and p-values may be found in Appendix Table 1. The various five-year bins for a are on the horizontal axis. The vertical axis represents $\alpha^{i,6}$ in percentage points. The solid line represents the (smoothed) point estimates. For a one percentage point unexpected increase in the federal funds rate, the point estimate ($\alpha^{6,a}$) represents the effect of an increase in the population share of group a by one percentage point on the growth rate of the variable of interest. The vertical bars depict the 95% confidence intervals for the estimated coefficients. These confidence intervals are estimated using the approach in Driscoll and Kraay

[8]Construction industry employment is not included for Delaware, Washington DC, Hawaii, Maryland, Nebraska, South Dakota, Tennessee.

[9]In 2015, the BEA began publishing quarterly state population estimates which are based on unpublished beginning of the month state population estimates from the Census Bureau.

(1998) which is robust to general forms of spatial and temporal correlation in the error terms.

The top panel presents results for personal income. The coefficients are positive for the young aged 20 to 39, and statistically significant at the one percent level for the 25-29 and 30-34 year olds. The coefficients are negative for the middle-aged 40 to 64-year olds, and statistically significant at the one percent level for all but the 40-44 and 60-64 year old group. The coefficients are small and insignificant for the old aged 65 and above. The positive coefficients for the young imply that if the fraction of the population in this age group increases, then an unanticipated *increase* in interest rates *raises* personal income after six quarters *relative* to the effect that monetary policy has on the US economy as a whole. Note that this does not mean that personal income raises, as the total effect of the money shock on state income is a combination of the first-order effect subsumed into the time effect and the interaction effect considered here. Since we typically expect an increase in interest rates to reduce personal income, a positive coefficient implies a weaker negative effect of monetary tightening. Similarly the negative coefficients for the middle-aged imply a stronger impact of monetary policy on personal income the larger is this age group.

The absolute size of the $\alpha^{6,a}$ have very little meaning. The magnitude of the coefficients depends on the size of our estimated monetary policy shocks, as well as the relationship between these shocks and the overall path for interest rates. We have normalized the size of the monetary policy shock to 1% for convenience. A one percentage point shock is not only very large relative to the shocks that we observe in our sample, but could be associated with a very different change in interest rates, especially after six quarters. While we cannot interpret the magnitude of the coefficients, we can get some sense of how large the effects are by comparing the $\alpha^{6,a}$ to the effects of monetary policy shocks on the US economy that we estimated in Appendix Section B. This comparison keeps the magnitude of the monetary shocks fixed. In Figure 1, we see that a one percent increase in the Federal Funds rate reduces personal income by almost five percentage points after six quarters. Now consider a one standard deviation change in the percentage of the population that is 30-34. This is an increase in p^{30-34} of 0.64 percentage points.[10] Given $\alpha^{6,30-34} = 3.13$, a one standard deviation increase in the share of 30-34 year olds reduces the impact

[10] This is the standard deviation of p^{30-34} after removing time effects. This is the relevant variation as discussed in Section 5.1.

Figure 2: Effects by age-bin (6-quarters after monetary shock)

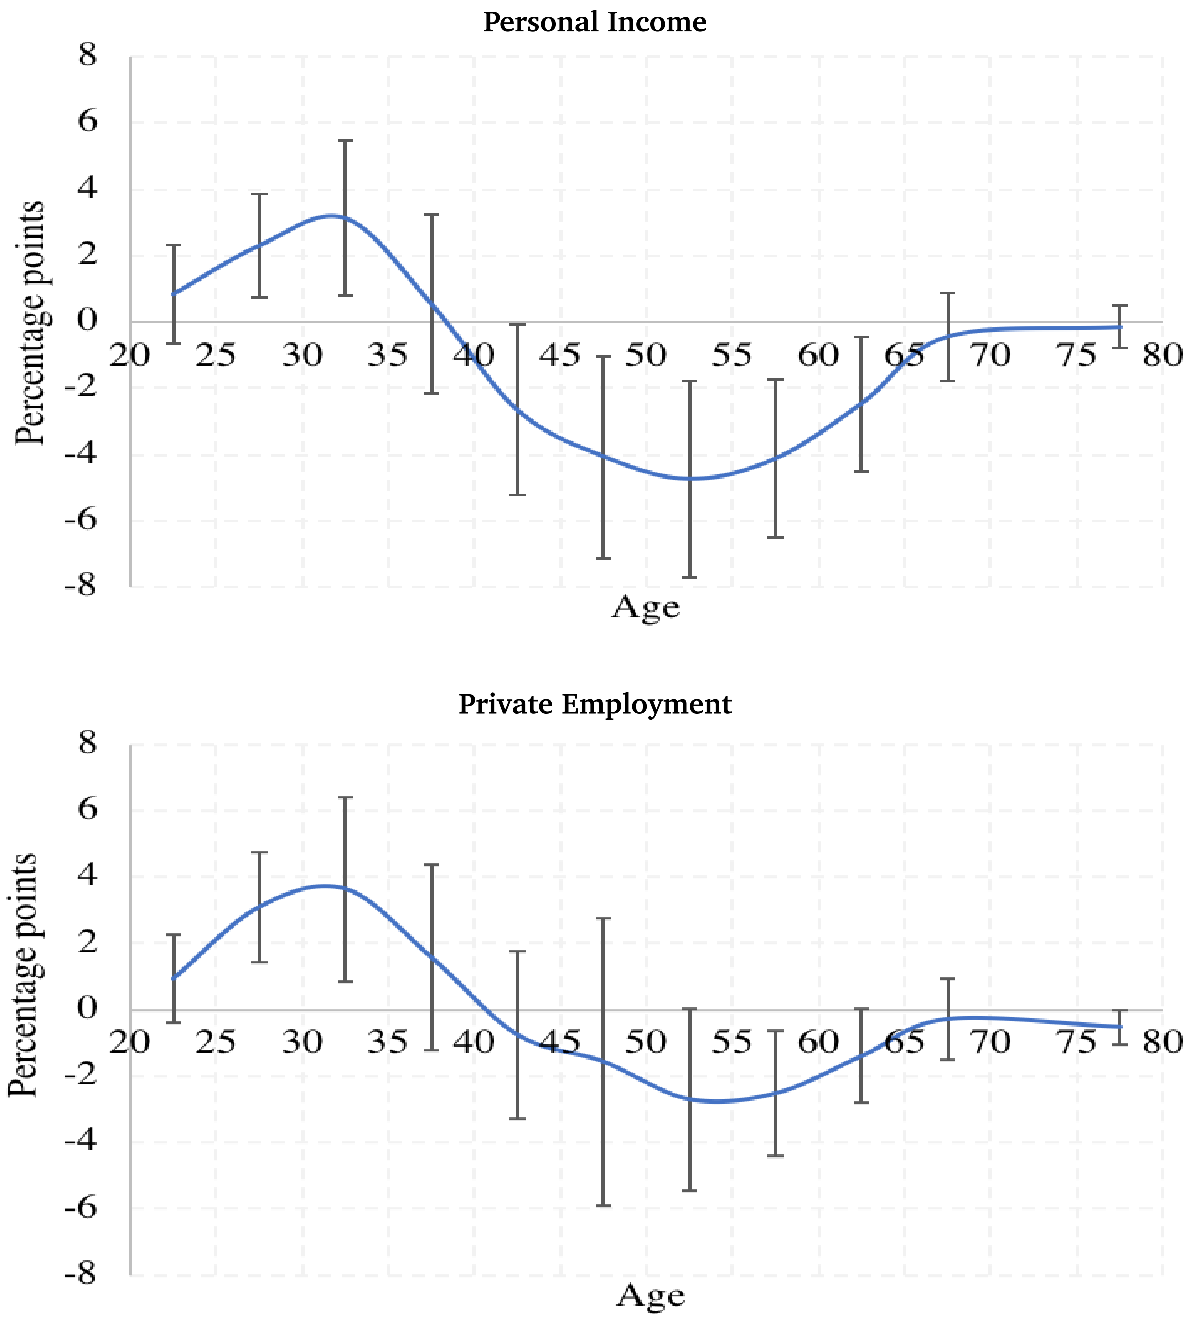

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

of a one percentage point increase in the Federal Funds rate by approximately two percentage points or 40% of the aggregate effect. Similarly, a one standard deviation increase in the share of 50-54 year olds equal to 0.39 percentage points together with $\alpha^{6,50-54} = -4.74$ implies that a one standard deviation increase in the share of 50-54 year olds increases the negative effects of a one percentage point increase in the Federal Funds rate by 1.85 percentage points. Again, approximately 40% of the aggregate effect. Hence the variation in the effects of a monetary policy shock induced by demography appear to be of a magnitude similar to the direct effects of monetary policy on the economy as a whole.

It appears that the sum of the coefficients in Figure 2 is approximately zero. Given that an increase in one age group must show up as a decrease in another age group, it is tempting to argue that the coefficients must sum to zero. This is not necessary, however, for three reasons. First we have omitted the 0-19 year old age group. Second, theoretically it is the variance weighted sum of coefficients that sums to zero.[11] Finally, since we run the regressions separately for each age group, variation in the coefficients on the other variables may affect this homogeneity property.

The bottom panel of Figure 2 illustrates the effects on private employment. The pattern closely mimics the results for personal income, but with some small differences in the turning points and a more moderate response for the middle age groups. The effects of monetary policy on private employment are dampened if there are more 20-34 year olds in the economy; the effects of monetary policy are strengthened if there are more 40-59 year olds in the economy; and are largely neutral if there is a change in the population over 60 years of age.

The pattern that we see in Figure 2 with positive coefficients for the young, negative coefficients for the middle-aged and small and statistically insignificant coefficients for the old is a pattern that we will see often in what follows. Generally, estimates will be stronger and statistically more significant for income than employment, reflecting the fact that both wages and employment respond in similar ways to the shock, so that the response of the product of the two components is stronger than the response of each component separately. Results for the young and the middle-aged are remarkably stable across alternative specifications. Results for the old are generally statistically

[11] If all of the age groups are included in a single regression then the OLS coefficient is $(X'X)^{-1}(X'Y)$. It is $X'Y$ that must sum to zero.

Figure 3: Impulse Response Functions of Personal Income

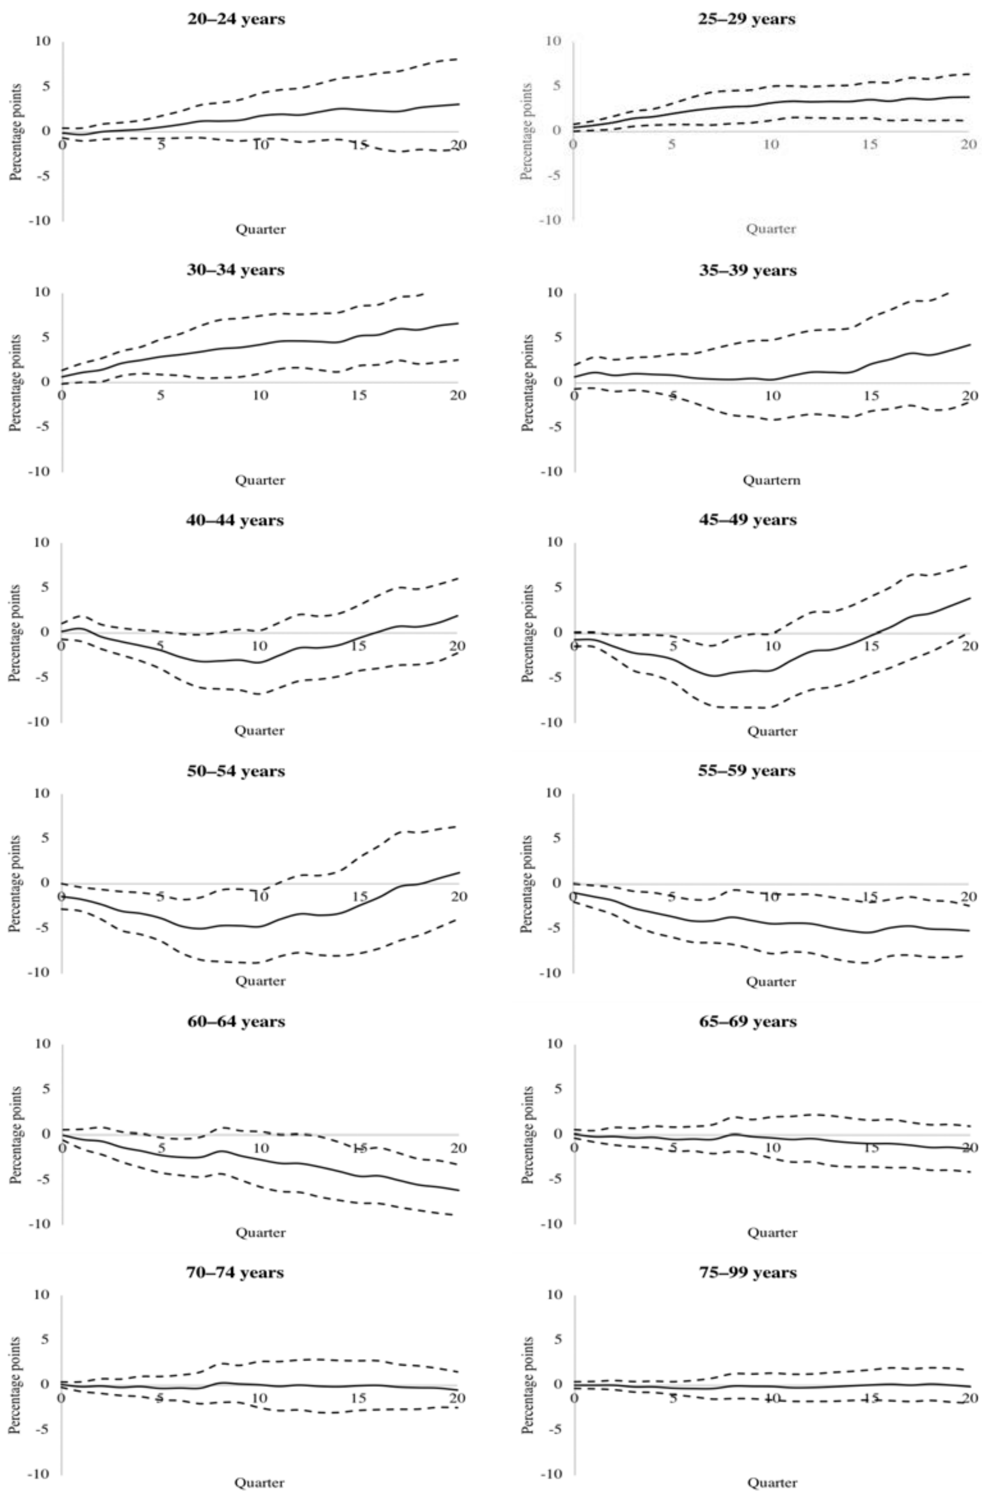

Each figure depicts the effects over time of a one percentage point increase in the share of the population in a specific age group on the responsiveness of the growth rate of state personal income to a positive one percentage point shock to the federal funds rate. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

insignificant but a bit less stable. There are a few specifications in which a greater share of old reduce the impact of monetary policy and a few specifications in which a greater share of the old increase the impact of monetary policy.

We presented the results for $i = 6$ because this is the horizon for which the effects of age appear strongest. We get qualitatively similar effects at other horizons over the first three to four years following a monetary shock. Figure 3 presents the $\alpha^{i,a}$ for personal income as functions of the horizon i for each age bin a. The effects of age tend to grow over the first six quarters, positive for the young and negative for the middle-aged. The maximal impact is at about six quarters. After six quarters the coefficients remain elevated for some age groups, especially the young, and die out for others, primarily those between 40 and 54. The effects are always small and statistically insignificant for the old.

4.2 Robustness

4.2.1 Monetary shocks

One obvious question is whether our results are driven by our choice of monetary policy shocks. Nakamura and Steinsson (2018) have raised questions regarding the interpretation of shocks identified though the high frequency movements of asset markets around policy announcements. These shocks may reflect unanticipated moments in policy that are orthogonal to current economic conditions, or they may reflect the communication of the Federal Reserve's information regarding current economic conditions.

We consider two alternative identification schemes. First, we consider the shocks that Romer and Romer (2004) construct from the minutes of Federal Open Market's Committee meetings and the Federal Reserve Board's Greenbook forecasts. These are extended through the end of 2007 by Wieland and Yang (2017). Second, we consider shocks estimated from a Structural VAR that includes US real GDP, the GDP deflator, the CRB spot commodity price index, and the quarterly average federal funds rate. The VAR is estimated over the 1960 to 2008 sample period and monetary policy shocks are identified via a Cholesky decomposition with the interest rate ordered last.

Figure 4 compares the effects on personal income and employment for the various shock specifications. In the top three rows, the sample is the same as in our baseline specification: 1990-2008. Given that different shocks have different scales, we normalize by the standard deviation of the shocks. The vertical axis is $\alpha^{i,a} * sd(\epsilon_t^m)$. For ease of comparison, all the figures are on the same scale. The first row reproduces the results from our baseline specification. Regardless of shocks, the effectiveness of monetary policy is dampened for the younger age groups, is strengthened for the middle age groups, and is about the same as the aggregate effect for the older age groups. While, there are slight differences in the boundaries between the three groups and differences in which bins are statistically significant, the basic shape and magnitude of the response is similar across the shock series.

With these alternative shock series, we can expand the sample period. The last two rows in Figure 4 shows the results for the 1980-2008 sample. The results for the VAR shocks become more significant and look much more like the baseline results. The results for the Romer-Romer shocks are largely unaffected by the sample period.

4.2.2 Other robustness checks

We tried a number of alternative specifications. These results are discussed in detail in Appendix C. Here we provide a brief overview of our findings.

- In the baseline specification we run separate regressions for each age group. We tried running a single regression with all of the age groups. To avoid a proliferation of regressors, we aggregated the age bins into three age groups: 20-390, 40-59, and 60+.
- We aggregated population and personal income by BEA regions.
- We dropped the five smallest states by population: Alaska, North Dakota, Vermont, Washington D.C., and Wyoming.
- We do not have a long enough sample to properly analyze subsample stability. We do, however, investigate whether our results are sensitive to altering the beginning and the end of the sample. First, we truncated all data at the end of 2006 to avoid any contamination with the housing market crash and the Great Recession. Second, we began our sample in 1995,

Figure 4: Alternative monetary shocks

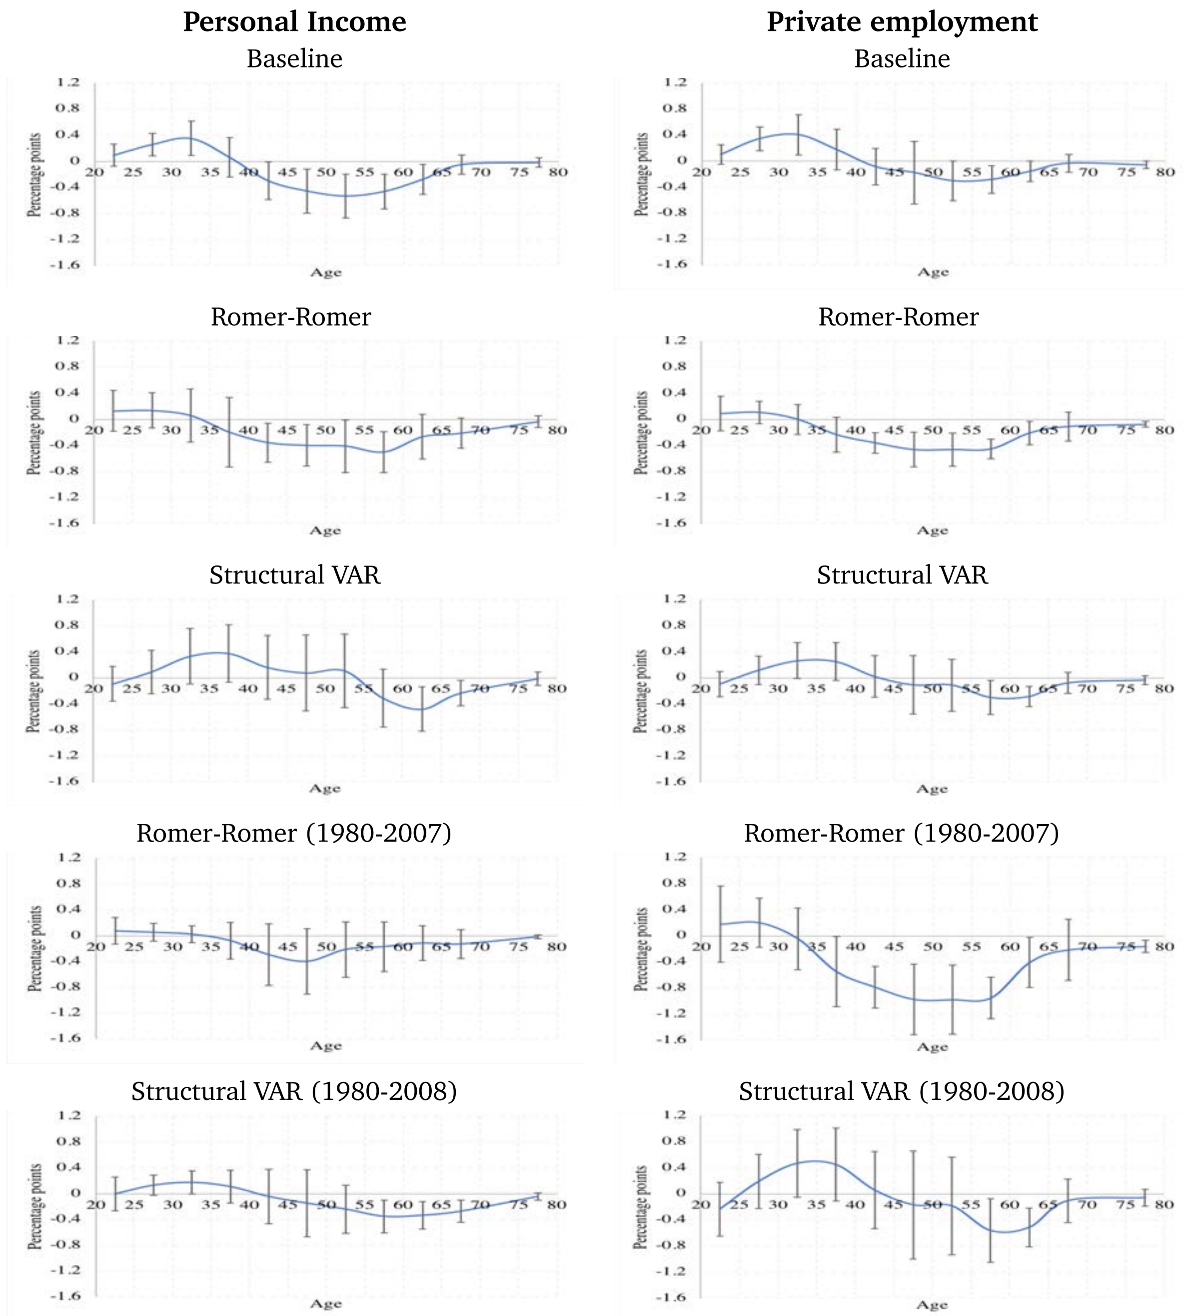

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors. The baseline and VAR specifications are estimated through 2008; Romer and Romer data ends in 2007.

after the FOMC began issuing press releases following every FOMC meeting. If anything, beginning in 1995 strengthened the economic and statistically significance of the estimates.

The bottom line as detailed in Appendix C is that the basic pattern seen in Figure 2 in which the young dampen the effects of monetary policy and the middle aged amplify the effects of monetary policy survives all of these modifications.

5 Inspecting the Mechanism

Given the robustness of the basic results, we turn our attention to potential mechanisms. We begin by addressing a couple of relatively uninteresting possibilities. First, it is possible that age structure is proxying for some other state characteristic, and that it is this characteristic that is driving the results. Second, it is possible that our results reflect migration between states rather than the expansion of opportunities within a state. While this second mechanism is not necessarily uninteresting, it would call for care in applying our results for US states to nation states. Migration is likely to be a more important margin of adjustment for US states where the borders are porous, than for nation states which may monitor and control entry and exit.

Next we ask whether the greater share of middle aged amplify the effects of monetary policy because the middle aged themselves respond more to the shock and so a greater number of middle aged increases the response, or because all age groups respond more to monetary policy shocks when the proportion of middle aged is larger. The evidence favors the second story. All age groups respond more. This points toward a general equilibrium mechanism by which the greater number of middle aged increases income for all age groups.

In the next section, we suggest such a mechanism. We show that proprietors income responds strongly to the age structure of the state and that greater share of middle aged increases the response of proprietors' income. This points to a mechanism by which profits drive business expansion which raises incomes of all groups in the state.

5.1 State Characteristics

One might be concerned that our results are driven by some correlation between the population structure of a state and some other state characteristics. We include time and state effects in the baseline specification (1). These effects absorb common trends and some state differences but not all. Since the money shock is common across states, the time effects control for trends in population growth that are common across states. To see this, consider adding a constant c_t to $p^a_{s,t}$ for all s at a given date t, then the interaction term becomes $(p^a_{s,t} + c_t)\varepsilon^m_t$, and $c_t\varepsilon^m_t$ is absorbed by the time effect. Given that the money shock varies over time, however, the state fixed effects do not absorb average population differences across states. To see this, consider adding a constant c_s to $p^a_{s,t}$ for a given state s for all t. The interaction term becomes $(p^a_{s,t} + c_s)\varepsilon^m_t$, and $c_s\varepsilon^m_t$ is not absorbed by either the time or state fixed effect. The relevant variation in (1) is therefore the population shares in a given age bin across states and time net of common time trends.

To illustrate this variation, Figure 5 represents the demographic composition of 50-54 year olds across states in 1990 and 2000. To construct the figure we first removed time fixed effects from the population share of this age group. We then sorted the residuals into five quintiles. Lighter colors represent states with higher residuals. For example, Hawaii shifted from the population share of 50-54 year olds being 0.3 percentage points below average in 1990 to 0.3 percentage points above average in 2000. Indiana, on the other hand, shifted from the population share being 0.2 percentage points above average in 1990 to 0.1 percentage points below average in 2000. Although there is substantial variation in the rankings across the two time periods, given that the state effects do not fully control for state characteristics, it is of interest to investigate whether our results are driven by some state characteristics that are correlated with the population shares.

Our first experiment is to include the interaction between the state fixed effect and the monetary policy shock as an additional regressor. This will soak up $c_s\varepsilon^m_t$ and control for level differences in population across states that might be correlated with fixed state characteristics. For example, Southern states may have more old and young but less industry. The results are presented in Appendix Section D and in particular Figure 15. The basic pattern is the same as in our baseline specification. There are a few notable differences. The response of personal income is amplified.

Figure 5: Changing demographic makeup of the U.S.: Ages 50-54

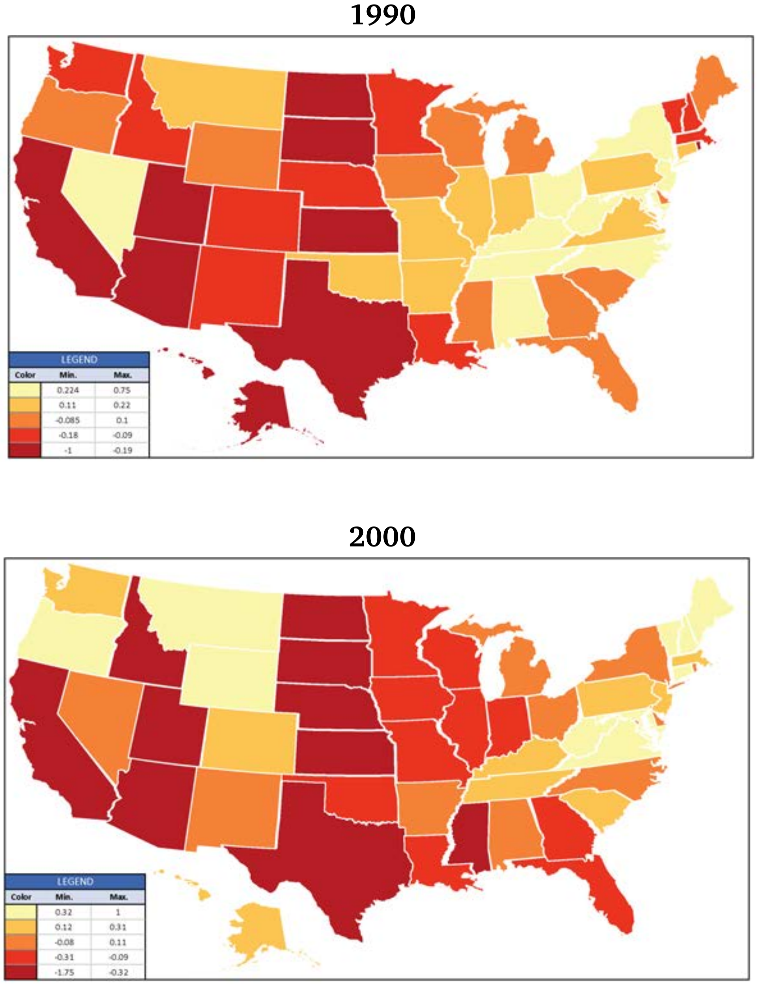

The figures represent the residuals of the share of population aged 50-54 years, after removing time fixed effects.

For both personal income and employment statistical significance improves for the 25-34 and 40-49 age groups, while worsening for those over 50. There is also weak evidence that a greater fraction of older people may reduce the effects of monetary policy on employment. Given that including these interaction terms does not affect the general pattern of the baseline results, in what follows we will report the results for our baseline specification that does not include the shock interacted with the state fixed effects and comment whenever including these interaction terms qualitatively alters the results.

Interacting the shock with state fixed effects controls for a correlation between average population in the state and average characteristics of the state. The possibility remains that some characteristics may change over time in ways correlated with the change in demographic structure. To investigate this issue we consider whether the population structure of a state is correlated with certain characteristics. We begin with the industrial composition of the state. We consider three industries which have been tied to monetary policy: manufacturing, financial services and construction. We consider manufacturing and construction because these industries tend to be more cyclical than other industries. A state with a large manufacturing or construction share would be expected to respond more strongly to a monetary policy shock. We consider financial services because of the effects of monetary policy through bank lending could explain a state's greater response.

We include the fraction of employment in manufacturing, construction and financial services both as additional regressors and interacted with the monetary policy shock. The addition of these additional regressors has no effect on the estimates of $\alpha^{6,a}$. We therefore present the results in Figure 16 in Appendix Section D. To understand why there is no effect, we consider the correlation between the share of employment in these industries and the proportion of the population between 40 and 64 years (both net of time effects). The correlation is -0.04 for manufacturing, -0.08 for financial activities, and -0.14 for construction.[12] All three correlations are small and negative. None of these sectors vary positively with the share of the middle aged as would be expected if these sectors were driving the results.

Industrial structure is not the only state characteristic that may confound our results. Cravino, Lan, and Levchenko (2018) argue that households with higher incomes tend to purchase goods with stickier prices. Since middle aged households tend to have high incomes and since stickier prices tend to be correlated with the response to monetary policy shocks, our results could reflect this mechanism. We therefore add state income interacted with the monetary policy shock as an additional regressor. Regression results appear in Figure 17 in Appendix Section D. Again we find no evidence that this mechanism is driving the results. The correlation between state personal income and the proportion of the population between 40 and 64 years of age (both net of time effects) is -0.11. Again the correlation is small and negative.

For completeness we considered other variables related to the demographic structure of the

[12]Using the level of industry employment instead of the fraction of industry employment to total private employment gives similar correlations.

economy, such as the percentage of employees who have college degrees, or are white, or are male, as well as some other state-level features of the economy that age might proxy for such as housing prices and homeownership rates. These results are presented in Figures 18 and 19 in Appendix Section D. We included each variable as well as the variable interacted with the monetary policy shock. None of these variables had a qualitative effect on the estimates of $\alpha^{6,a}$.

While we cannot conclude with certainty that the results of Section 4.1 are not due to the correlation of population structure with some other state characteristic, we find no evidence that this is the case.

5.2 Migration

One of the margins by which states adjust to shocks is through migration (Blanchard and Katz, 1992). This raises the question of whether the responses in Figure 2 result from hiring within the state or from migration between states. It also raises the question as to whether our results for US states can be applied to nations, since migration between states is likely easier than migration between nations.

To investigate this question we consider the responses of population and unemployment to monetary policy shocks interacted with the age bins. The idea is that migration should be reflected in the response of population, whereas hiring from within the state should show up in unemployment.

In the left panel of Figure 6, we take X to be the log change in population over six quarters, $\ln P_{s,t+6} - \ln P_{s,t-1}$, where $P_{s,t}$ is the population of state s at data t. The response of population looks very different than the response of employment in Figure 2. Whereas the response of employment and personal income partitioned the population into three three distinct age groups, the response of population partitions the population into only two. Population is more responsive in states with a greater share of the population younger than 50 and less responsive the greater the share over 50. This likely reflects the fact that the young are more mobile than the old.[13] The response of population is strongest for those 35-45 and those older than 60. The response of income and employment is generally the weakest for these age bins. Moreover, the response of population is

[13] As might be expected the magnitude of the response is muted if we drop the five States with the smallest population. The basic shape remains the same.

less than half the response of employment. We therefore conclude that migration, while potentially important, is not the main factor behind our results.

Figure 6: Effects of Migration

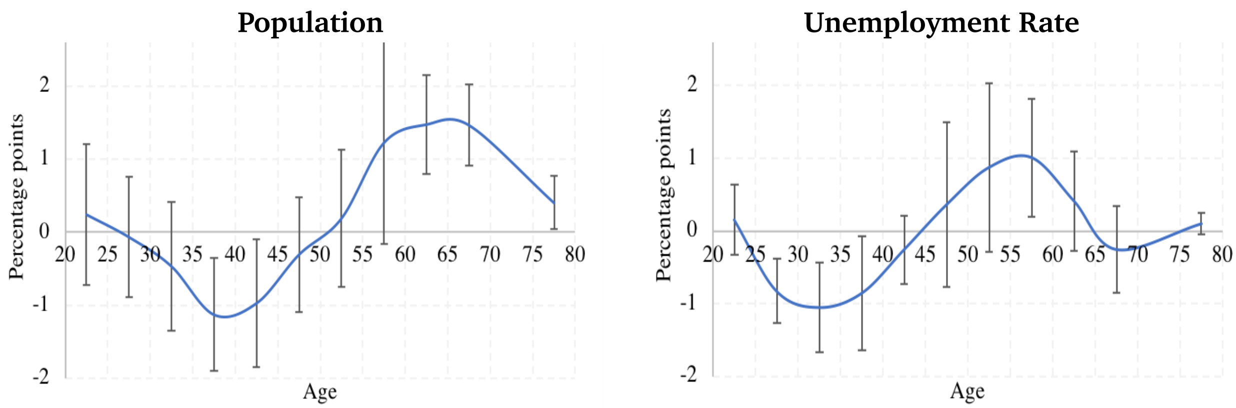

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

In the right panel, we take the dependent variable to be the change in the unemployment rate, $U_{s,t+6} - U_{s,t-1}$. This figure looks very similar to Figure 2. The sign of the response is flipped as an increase in unemployment tends to coincide with a decline in employment. Together these two figures point to the importance of internal adjustment mechanisms and increase our confidence that conclusions that we draw for the US states should apply to nation states.

5.3 Aggregate or Individual Channels

Our next experiment is to replace X with $X^{\bar{a}}$, the income of age group $\bar{a}$, to investigate whether the mechanism works at the level of the individual or in the aggregate. Let $\alpha^{i,a}(\bar{a})$ denote the response of $X^{\bar{a}}$ to $\varepsilon^m p^a$. If $\alpha^{i,a}(\bar{a})$ is the same for all a, then the response of each age group $\bar{a}$ is independent of the age distribution in the state. In this case, the mechanism likely operates at the individual level, and the aggregate response reflects the cumulated individual responses. If, on the other hand, $\alpha^{i,a}(\bar{a})$ is the same for all $\bar{a}$, then all age groups $\bar{a}$ respond similarly to an increase in the proportion of the population in bin a. In this case the mechanism likely operates at the aggregate level.

We use data from the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Program. The LEHD provides average monthly earnings, the number of jobs, and payroll for employees. The publicly available data is available for the following pre-defined age bins, 19-21 years, 22-24 years, and then 10-year bins from 25-34 years to 55-64 years, followed by one last 65-99 years age bin. We constuct the income age groups $\bar{a}$ to match the LEHD age-bins.

We find strong evidence for aggregate mechanisms. Figure 7 depicts $\alpha^{6,a}(\bar{a})$ for different choices of $\bar{a}$ and for different choices of X. $X^{\bar{a}}$ is either average earnings per worker, the number of jobs, or total payroll in bin $\bar{a}$. The overall picture remains similar to section 4.1 above. The response of average earnings, jobs, and payrolls are independent of one's own age group and depend primarily on the age structure of the state in which one lives. It is not so much being middle aged as living in a state with lots of middle aged residents that explains the response of income to a monetary policy shock.[14]

As further evidence that the mechanism operates through an aggregate channel, Figure 7 shows that the mechanism appears to work both through the extensive margin, the number of jobs, and the intensive margin, earnings per worker. The combined effect on total payroll is the product of the two. This combined effect is stronger and statistically more significant.

6 Middle-Aged Entrepreneurs and Monetary Policy

In this section, we suggest a mechanism by which the greater number of middle aged amplify the impact of monetary policy on incomes and employment of all age groups. Our mechanism is motivated by the observation that entrepreneurs, especially successful entrepreneurs, tend to be middle aged (Selgado, 2018, Azoulay et al., 2019). Suppose that middle aged residents, due to their experience and resources, have a comparative advantage in entrepreneurship. Then states with more middle aged residents will likely have more entrepreneurs on the margin of entry. Business formation will then respond more aggressively to a reduction in interest rates and income and employment will follow. These first round effects will be followed by knock-on effects. The relatively larger expansion in states with more potential entrepreneurs will tend to increase business

[14]When we include the interaction between the monetary shock and the state fixed effect as an additional control, Figure 7 is largely unchanged except that $\alpha^{i,a}$ is often positive and significant for the old, age 65 and over.

Figure 7: Which Age Group's Income Matters Most?

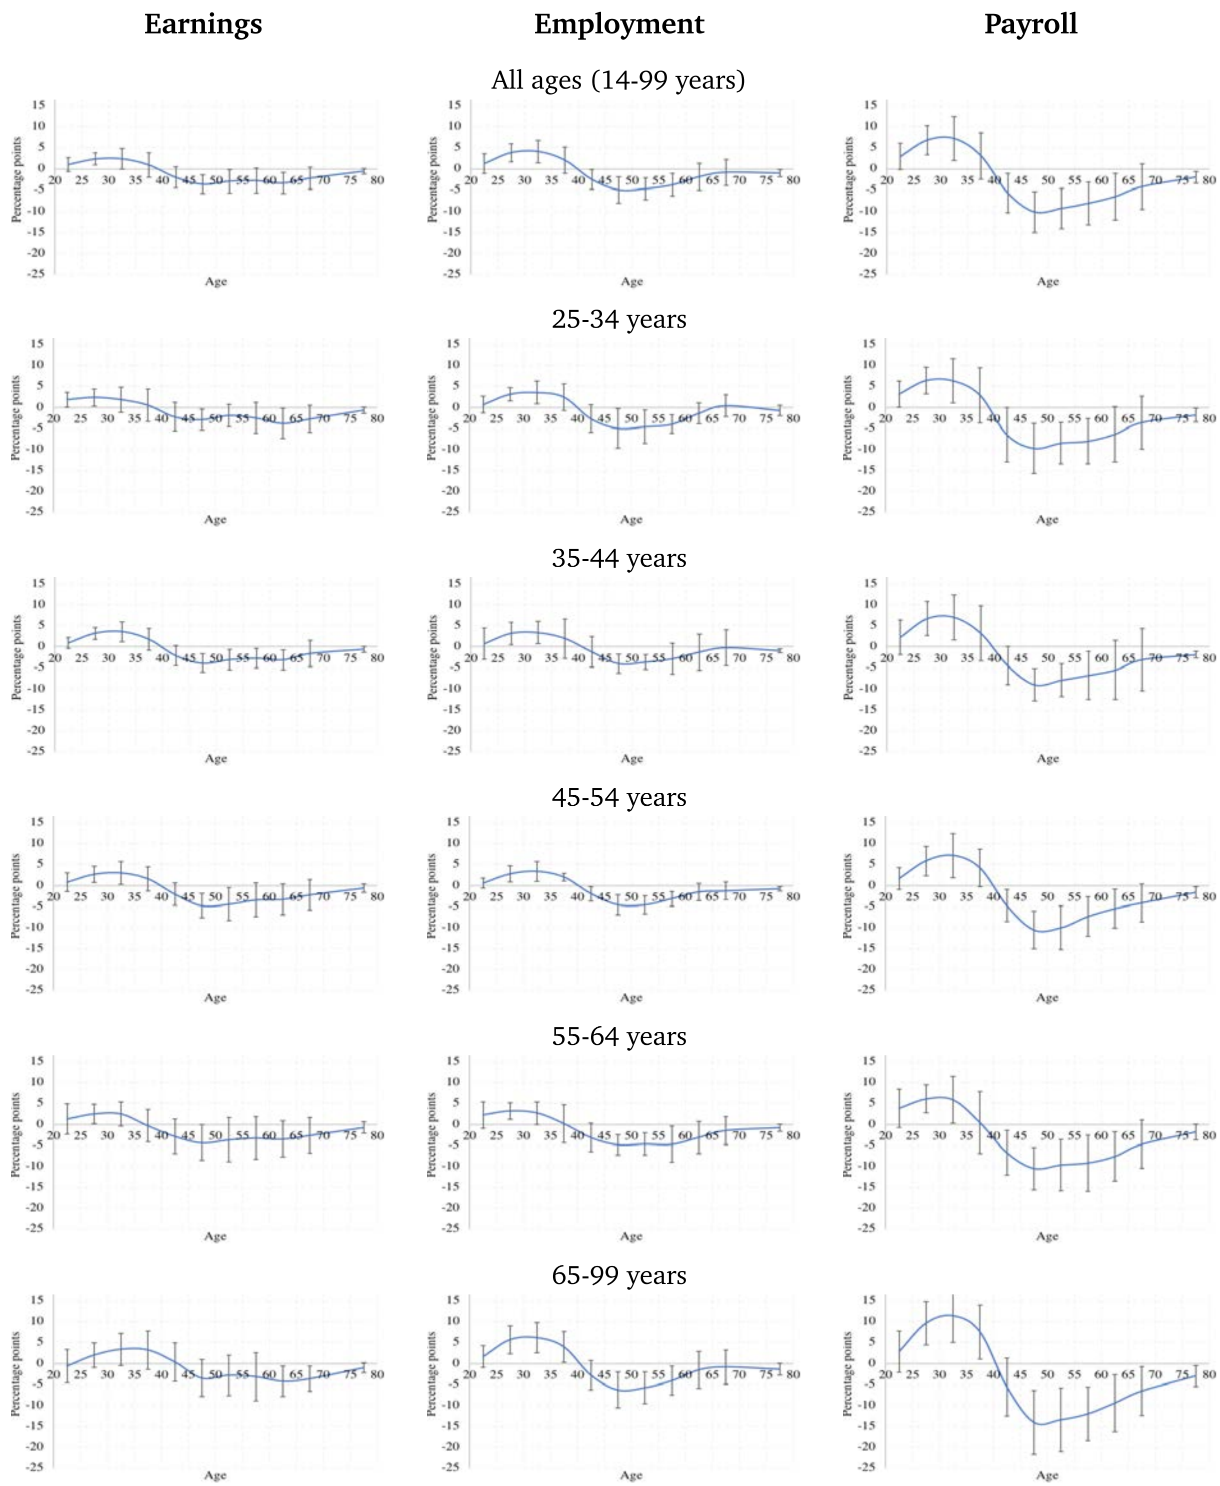

Each graph represents the response of earnings, employment, or payroll of all individuals in a specific age-group (listed in the row titles) to a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a (listed on the horizontal axis of each figure). The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

income. The rise in business income will further increase the incentive to invest and will provide entrepreneurs with the resources to take advantage of investment opportunities and expand employment. The result of all this investment will be a general rise in output and incomes that benefit all age groups spurred by the entrepreneurial activities of the middle aged. Moreover, the increased incentive to invest will tend to cause the middle-aged to substitute out of consumption, helping to reconcile our results with the earlier literature that finds it is the consumption of the young that responds more to monetary shocks.

We find evidence for certain aspects of this story. We find that establishment births respond more strongly to monetary policy the greater the share of middle aged (See Figure 8, left panel). Establishment births slightly lead income and employment. They peak five quarters after the shock, one quarter prior to the peak in income. Interestingly, we find no consistent response of establishment deaths. The response of deaths is generally insignificant and the estimates are unstable over time (See Figure 8, right panel).[15] In some sense this is not surprising. There are two effects of a monetary expansion. The first is to make existing firms more profitable. This effect tends to reduce establishment deaths. The second is to increase the entry of new firms run by entrepreneurs with marginal ability. This effect tends to increase the rate of business failure. The net effect on establishment deaths is therefore ambiguous.

We also find that business profits respond more the greater the share of middle aged. Figure 9 illustrates the response of different types of income to our interacted shock. The response of wages and salaries and the response of dividend, interest, and rent look very similar in shape and magnitude to our baseline specification. The main difference is in the response of proprietors' income. Here the negative response of the middle-aged is much more pronounced, although imprecisely estimated.[16]

Consistent with the story that increased profits further drive business expansion, we find that the share of small business employment rises more rapidly the larger the fraction of the middle

[15]When we include the interaction between the monetary shock and the state fixed effect as an additional control, Figure 8 is largely unchanged except that in the case of establishment deaths side the coefficient for the 40-44 year-old bin becomes positive and significant.

[16]The response of proprietors income in Figure 9 is large but imprecisely measured. If we include the monteary shock interacted with State fixed effects the respone of proprietors income remains large and is statistically significant. Results available from the authors upon request.

Figure 8: Business Formation

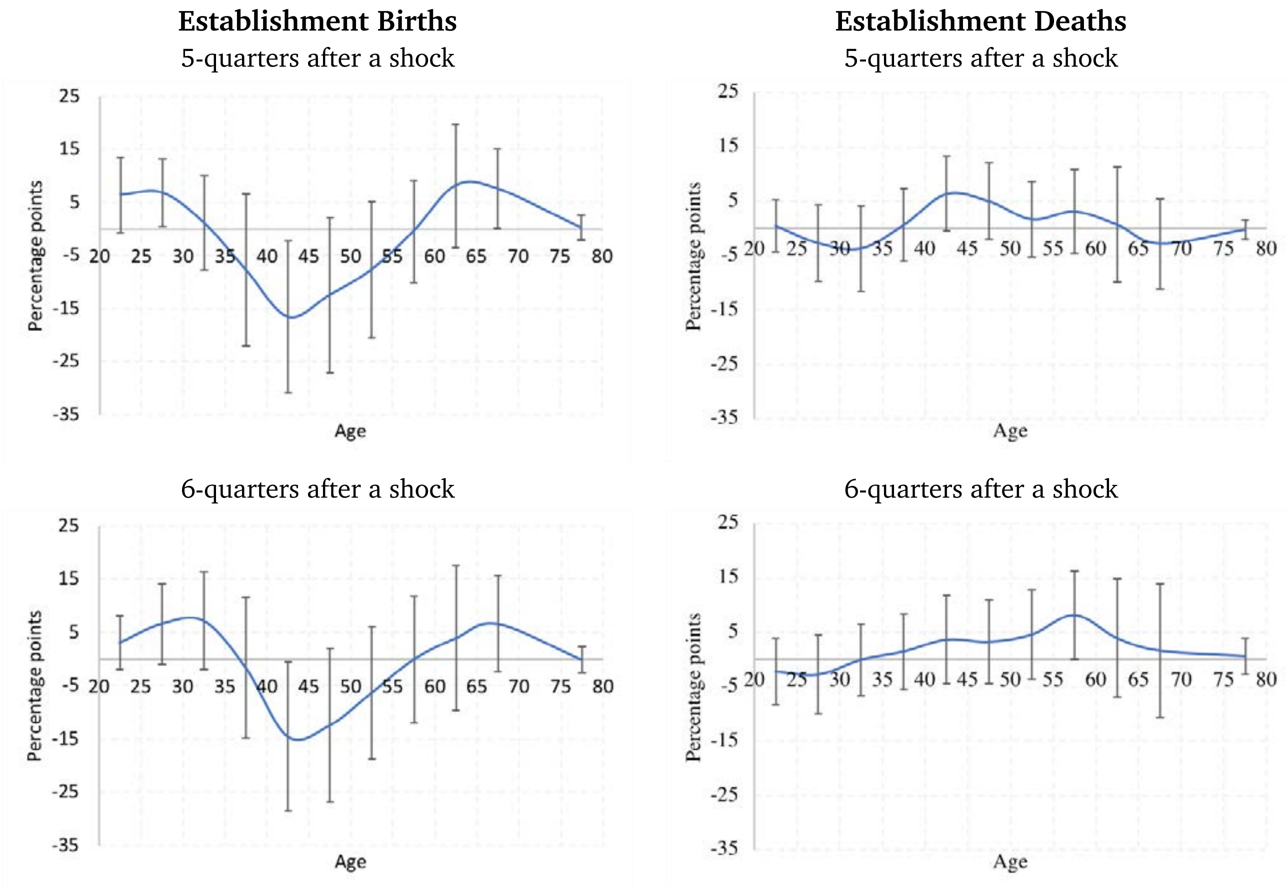

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

aged (See Figure 10). This effect is reminiscent of the effects found in Gertler and Gilchrist (1994) in which monetary policy has a greater impact on small firms. Gertler and Gilchrist interpret this as evidence of financial frictions. In Figures 20 and 21 Appendix Section E, we show the response of firms by size and age. While firms of all types show the familiar pattern in which the impact of monetary policy is lessened for the young and enhanced for the old, this pattern is more pronounced in the young and small firms.[17]

One might worry that the share for the middle-aged is simply proxying for firm size. It is true

[17] In Figure 10 we define a small firm as any firm with fewer than 500 employees. The pattern illustrated in the figure does not depend on the cutoff used to define a small firm. A similar pattern is visible if we take the cutoff to be 50 or 250 employees.

Figure 9: Response of Different Types of Income

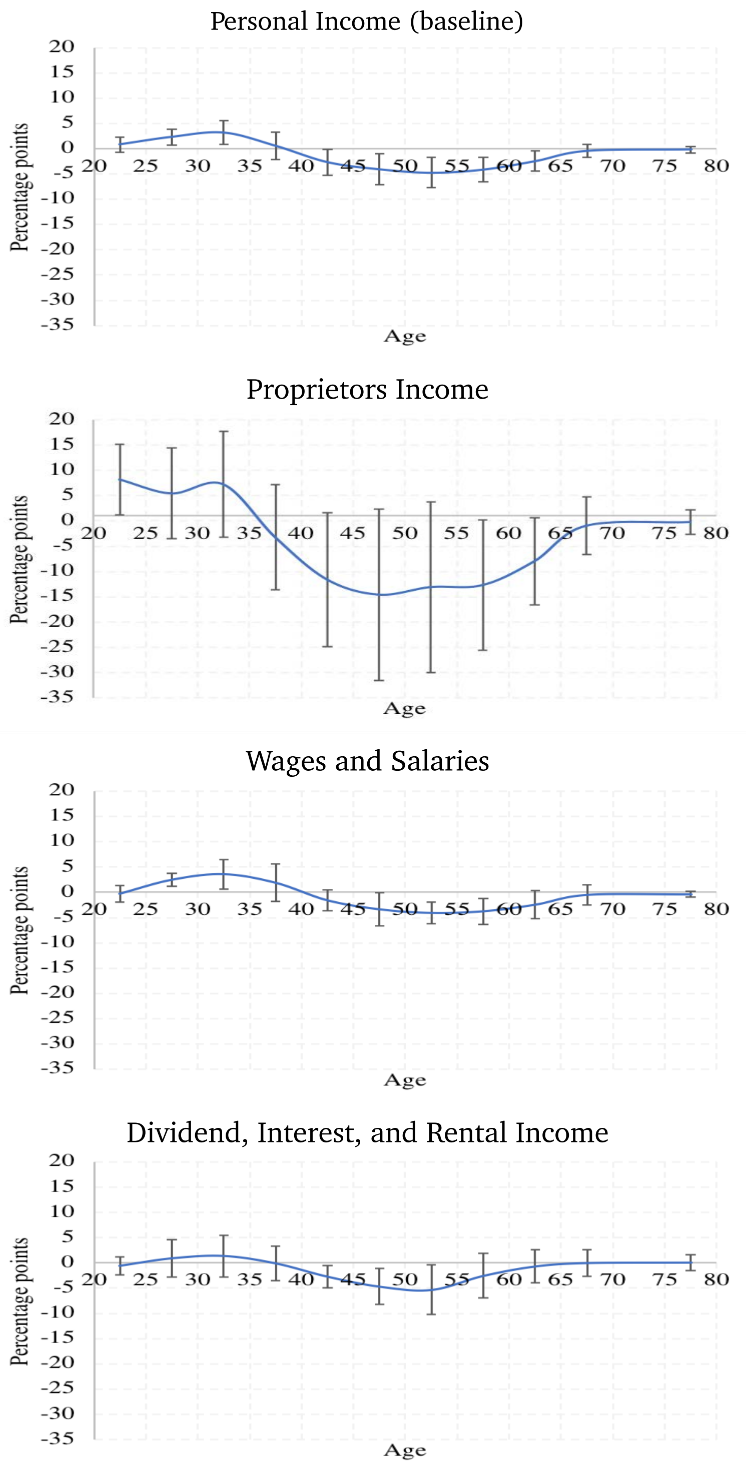

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 10: Response of employment share of small firms to all firms

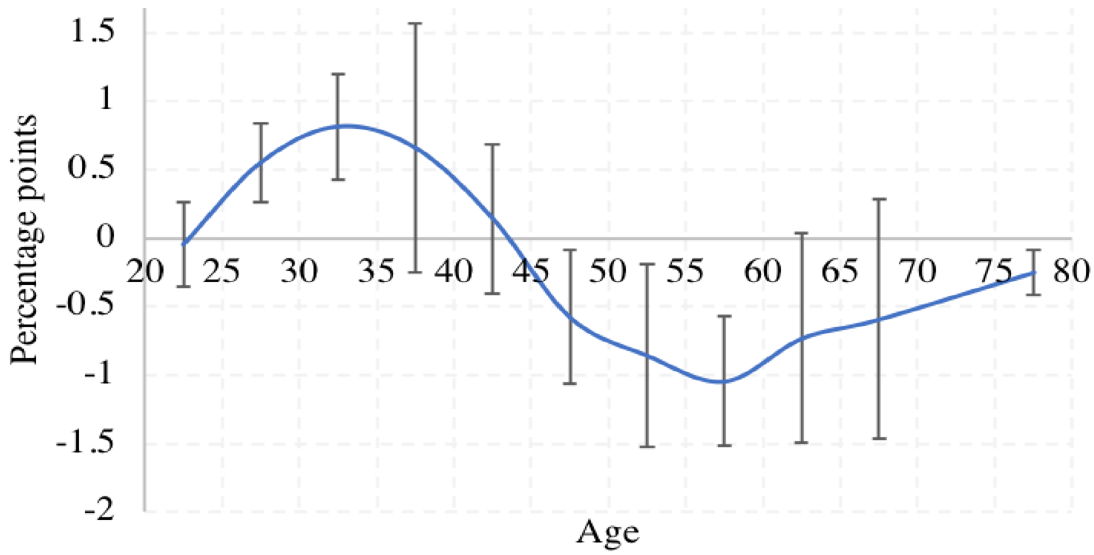

Small firms are defined as firms with 0-499 employees. The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

that the share of the middle-aged is correlated with the share of small business employment, as would be expected if the middle aged tended to be the ones to start businesses. The correlation between the share of the population 40-64 and the share of employment in businesses with fewer than 500 employees is 0.32 (after removing time effects). It does not appear, however, that the middle-aged population is simply a proxy for small business. Appendix Figure 22 shows that the pattern seen in Figure 2 survives the inclusion of the small business share interacted with the monetary policy shock. The evidence is therefore consistent with a story in which a monetary expansion prompts entrepreneurs to start and expand small businesses, and this expansion leads to a rise in the incomes of all age groups.

7 Conclusion

We consider the effects of monetary policy shocks on income and employment in U.S. states. We find that there are three distinct age groups in each state. The young, who are under 35 years of age, tend to dampen the effectiveness of monetary policy on personal income and employment. The middle aged, who are between 40 and 64 years, tend to amplify the impact of monetary policy. The old, who are over 65 years of age, do not affect the impact of monetary policy one way or another. We find that these effects work through the aggregate, not the individual: what matters for the income of a given age group is not the age group itself, but the age composition of the state. We find that these results cannot be explained by the industry composition of a state. We propose a mechanism by which a greater share of middle aged residents increases the elasticity of business formation in response to monetary policy.

Our results are complementary to the findings of Sterk and Tenreyro (2018), Wong (2018), and Cloyne et al. (2018). These authors study the effect of monetary policy on the consumption of different age groups. We study how the effect of monetary policy on aggregate income varies with the age structure of the economy. An interesting direction for future research would be to construct a model consistent with both sets of facts. In the model, liquidity constraints would cause the consumption of the young to respond strongly to monetary policy shocks. The middle aged would respond, not by consuming more but by investing, starting business, and hiring, all activities which lead the economy to expand and increase the incomes of all groups.

References

[1] Auclert, Adrien (2019), “Monetary Policy and the Redistribution Channel,” *American Economic Review*, forthcoming.

[2] Azoulay, Pierre, Benjamin Jones, Daniel Kim and Javier Miranda (2019), “Age and High-Growth Entrepreneurship,” *American Economic Review: Insights*, forthcoming.

[3] Basso, Henrique, and Omar Rachedi (2018), “The Young, the Old, and the Government: Demographics and Fiscal Multipliers,” Bank of Spain working paper.

[4] Berg, Kimberly, Chadwick Curtis, Steven Lugauer, and Nelson Mark (2019), “Demographis and Monetary Policy Shocks,” working paper.

[5] Blanchard, Olivier, Lawrance Katz (1992), “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1 :1-75.

[6] Chodorow-Reich, Gabriel (2018), “Geographic Cross-Sectional Fiscal Multipliers: What Have We Learned?” *American Economic Journal: Economic Policy*, forthcoming.

[7] Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio (2016), “Demographic and Real Interest Rates: Inspecting the Mechanism,” *European Economic Review* 88, 208-226.

[8] Carvino, Javier, Ting Lan, and Andrei Levchenko (2018), “Price Stickiness along the Income Distribution and the Effects of Monetary Policy,” *Journal of Monetary Economics*, forthcoming.

[9] Cloyne, James, Clodomiro Ferreira, and Paolo Surico (2018), “Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism,” Banco de España Working Paper 1813.

[10] Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng and John Silvia (2017), “Innocent Bystanders? Monetary Policy and Inequality in the U.S.,” *Journal of Monetary Economics* 88, 70-88.

[11] Doepke, Matthias, and Martin Schneider (2006), “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 114 (6), 1069-1097.

[12] Driscoll, John and Aart Kraay (1998), “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data,” *Review of Economics and Statistics* 80, 549-560.

[13] Eggertsson, Gauti, Neil Mehrotra, and Jacob Robbins (2017), “A Model of Secular Stagnation: Theory and Quantitative Evidence,” NBER Working Paper No. 23093.

[14] Freyer, James, “Demographics and Productivity,” *The Review of Economics and Statistics*, February 2007, 89, 100-109

[15] Fujiwara, Ippei, and Yuki Teranishi (2008), “A Dynamic New Keynesian Life-Cycle Model: Societal Aging, Demographics, and Monetary Policy,” *Journal of Economic Dynamics & Control* 32, 2398-2427.

[16] Gertler, Mark (1999), "Government Debt and Social Security in a Life-cycle Model." *Carnegie-Rochester Conference on Public Policy* 50, 61-110.

[17] Gertler, Mark, and Simon Gilchrist (1994), "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics* 109, 309-340.

[18] Gertler, Mark and Peter Karadi (2015), "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics* 7, 44-76.

[19] Gagnon, Etienne, Benjamin Johannsen, and David Lopez-Salido (2016), "Understanding the New Normal: The Role of Demographics," Finance and Economics Discussion Series 2016-080. Board of Governors of the Federal Reserve System (U.S.).

[20] Gürkaynak, Refet, Brian Sack, and Eric Swanson (2005), "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," International Journal of Central Banking.

[21] Gorodnichenko, Yuriy, and Michael Weber (2015), "Are sticky prices costly? Evidence from the stock market," *The American Economic Review*, 106, 165-199.

[22] Gourinchas, Pierre Olivier and Jonathan Parker (2002), "Consumption over the Life-Cycle," *Econometrica* 70, 47-89.

[23] Imam, Patrick (2015), "Shock from Graying: Is the Demographic Shift Weakening Monetary Policy Effectiveness," *International Journal of Finance & Economics* 20, 138–154.

[24] Kantur, Zeynep (2013), "Aging and Monetary Policy," working paper.

[25] Kaplan, Greg, Benjamin Moll, and Gianluca Violante (2018), "Monetary Policy According to HANK," *American Economic Review*, 108, 697?743.

[26] Kara, Engin, and Leopold von Thadden (2016), "Interest Rate Effects of Demographic Changes in a New Keynesian Life-Cycle Framework," *Macroeconomic Dynamics* 20, 120–164.

[27] Kuttner, Kenneth N. (2001), "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 47 (3), 523-544.

[28] Nakamura, Emi, and Jon Steinsson (2014), "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review* 104, 753-92.

[29] Nakamura, Emi, and Jon Steinsson (2018), "High Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 133(3), 1283-1330.

[30] Ottonello, Pablo, and Thomas Winberry (2018), "Financial Heterogeneity and the Investment Channel of Monetary Policy," NBER working paper No. 24221.

[31] Ramey, Valery (2016), "Macroeconomic Shocks and their Propagation," in John Taylor and Harald Uhlig eds., *Handbook of Macroeconomics*, Elsevier.

[32] Romer, Christina, and David Romer (2004), "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055-1084.

[33] Rudebusch, Glenn (1998), "Do Measures of Monetary Policy in a VAR Make Sense?," *Interna-*

tional Economic Review, 39(4), 907-931.

[34] Selgado, Sergio (2018), "Technological Change and Entrepreneurship," working paper.

[35] Shoag, Daniel (2010), "The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plan Returns." working paper.

[36] Söderström, Ulf (2001), "Predicting monetary policy with federal funds futures prices," *Journal of Futures Markets*, 21(4), 377-391.

[37] Sterk, Vincent and Silvana Tenreyo (2018), "The Transmission of Monetary Policy through Redistribution and Durable Purchases," *Journal of Monetary Economics*, forthcoming.

[38] Thapar, Aditi (2008), "Using Private Forecasts to Estimate the Effects of Monetary Policy," *Journal of Monetary Economics*, 55, 806-824.

[39] Wong, Arlene (2018), "Transmission of Monetary Policy to Consumption and Population Aging," working paper.

[40] Wieland, Johannes, and Mu-Jeung Yang (2016), "Financial Dampening," NBER working paper No. 22141.

Table 1: Baseline Results

	Personal Income			Private Employment		
	Driscoll-Kray			Driscoll-Kray		
	Coefficient	Standard Error	P-Value	Coefficient	Standard Error	P-Value
20-24 years	0.83	0.74	0.265	0.93	0.66	0.169
25-29 years	2.30	0.77	0.004	3.08	0.82	0.000
30-34 years	3.13	1.17	0.010	3.63	1.38	0.011
35-39 years	0.54	1.35	0.689	1.57	1.39	0.266
40-44 years	-2.66	1.28	0.043	-0.76	1.26	0.549
45-49 years	-4.06	1.52	0.010	-1.57	2.15	0.468
50-54 years	-4.74	1.48	0.002	-2.70	1.36	0.053
55-59 years	-4.12	1.18	0.001	-2.52	0.94	0.010
60-64 years	-2.46	1.01	0.018	-1.40	0.70	0.052
65-69 years	-0.45	0.67	0.506	-0.29	0.61	0.642
70-99 years	-0.15	0.32	0.628	-0.53	0.25	0.044

The first column contains the age groups, a. The coefficients represent the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a.

Appendix

A Baseline Coefficient Estimates

Table 1 presents the coefficient estimates behind Figure 2.

B Monetary Policy Shocks

This appendix explains the construction of our monetary policy shock series. This identification method involves using intra-day data on federal funds futures contracts around the time of monetary policy announcements. Monetary policy shocks are estimated as the change in the federal funds rate implied by the current month federal funds futures contract in a relatively short window of time around the FOMC announcement. By considering a sufficiently narrow window around the

FOMC announcement, one can be reasonably certain that no other news caused the change in the futures rates.

We construct our shock series in two steps. First, we obtain raw shocks from the policy announcements. Second, we aggregate these shocks to a quarterly frequency.

The raw shocks to monetary policy $\left(\epsilon_{\tau}^{m,hf}\right)$ are defined as follows:

$$\epsilon_{\tau}^{m,hf} = \frac{\overline{d}_{\tau}}{\overline{d}_{\tau} - d\left(\tau\right)} \times \left(i_{\tau+\Delta_1} - i_{\tau-\Delta_0}\right), \tag{3}$$

Here τ is the date and time of a FOMC announcement, $i_{\tau+\Delta_1}$ is the federal funds rate implied by the current month federal funds futures contract Δ_1 minutes after the FOMC announcement, while $i_{\tau-\Delta_0}$ is the rate implied by futures markets Δ_0 minutes before the announcement. We consider a one hour window around the announcements, from fifteen minutes prior to the announcement to forty-five minutes after the announcement. Positive values of $\epsilon_{\tau}^{m,hf}$ here imply unexpected increases (monetary contractions) in the federal funds rate and vice-versa. The coefficient multiplying the change in the interest rate, accounts for the fact that the settlement price of federal funds futures contracts is based on the realized average effective federal funds rate for the calendar month of the contract and the monetary policy shock implicit in the announcement only affects the federal funds rate over the remainder of the month. Here $\overline{d}$ is the maximum number of days in the month and $d\left(\tau\right)$ is the day of the month of the FOMC announcement.

Following Ottonello and Winberry (2018), we aggregate the high-frequency shocks to a quarterly frequency using a weighted moving average of the raw shocks,

$$\epsilon_t^m = \sum_{\tau \in t} \frac{\overline{d}_t - d\left(t,\tau\right)}{\overline{d}_t} \times \epsilon_{\tau}^{m,hf} + \sum_{\tau \in t-1} \frac{d\left(t-1,\tau\right)}{\overline{d}_t} \times \epsilon_{\tau}^{m,hf}. \tag{4}$$

Here t measures time in quarters and the ϵ_t^m are the quarterly aggregated monetary policy shocks that we use in our analysis. $\overline{d_t}$ is the number days in quarter t, and $d(t,\tau)$ is the day in quarter t of announcement τ. This method of aggregating shocks implies that an announcement made on the last day of a quarter will have a very small weight on the current quarter, but a relatively large impact on the next quarter. It is based on the idea that shock at the very end of quarter $t-1$ looks much more like a shock at the very beginning of quarter t than a shock at the very beginning of quarter $t-1$.

Trading in federal funds futures markets began in October 1988. Following Gürkaynak, Sack, and Swanson (2005), we begin our sample in 1990 and end in 2008 when interest rates hit the zero lower bound. We use the federal funds futures shocks from Gürkaynak, et al. (2005) and Gorodnichenko and Weber (2015). Over our baseline sample period of 1990-2008, there were eight scheduled Federal Open Markets Committee (FOMC) meetings in every calendar year. As discussed in detail in Gürkaynak, et al. (2005), the current practice of issuing a press release after every FOMC meeting began in 1994. Over the 1990-94 period, financial markets inferred the size and direction of the target federal funds rate based on open market operations conducted at 11:30am on the first business day after the FOMC meeting. In addition, on some occasions, a press release regarding discount rate changes was issued prior to an FOMC meeting. In these situations, financial market participants correctly inferred a change in the federal funds rate as well.

Figure 11 shows the actual changes in the target federal funds rate as well as the exogenous shocks used in our baseline representation in this paper. The dashed line depicts the actual change in the target federal funds rate while the solid line represents the high-frequency shocks as described above. The identified shocks are small but not insignificant. The average high-frequency shock over the 1990-2008 sample is about zero with a standard deviation of 11 basis points. Approximately 40% of shocks were contractionary while 60% are expansionary, and there are 70% more shocks than actual changes in the federal funds rate.

C Robustness

We consider a number of robustness checks to our baseline specification in this section.

C.1 Wide Age Bins

In our baseline specification, we estimate separate regressions for each five-year age bin. This allows the coefficients on the fixed effects and the other regressors to change with each specification. It also does not take into account that an increase in the proportion of the population in one

Figure 11: Federal funds rate changes and shocks

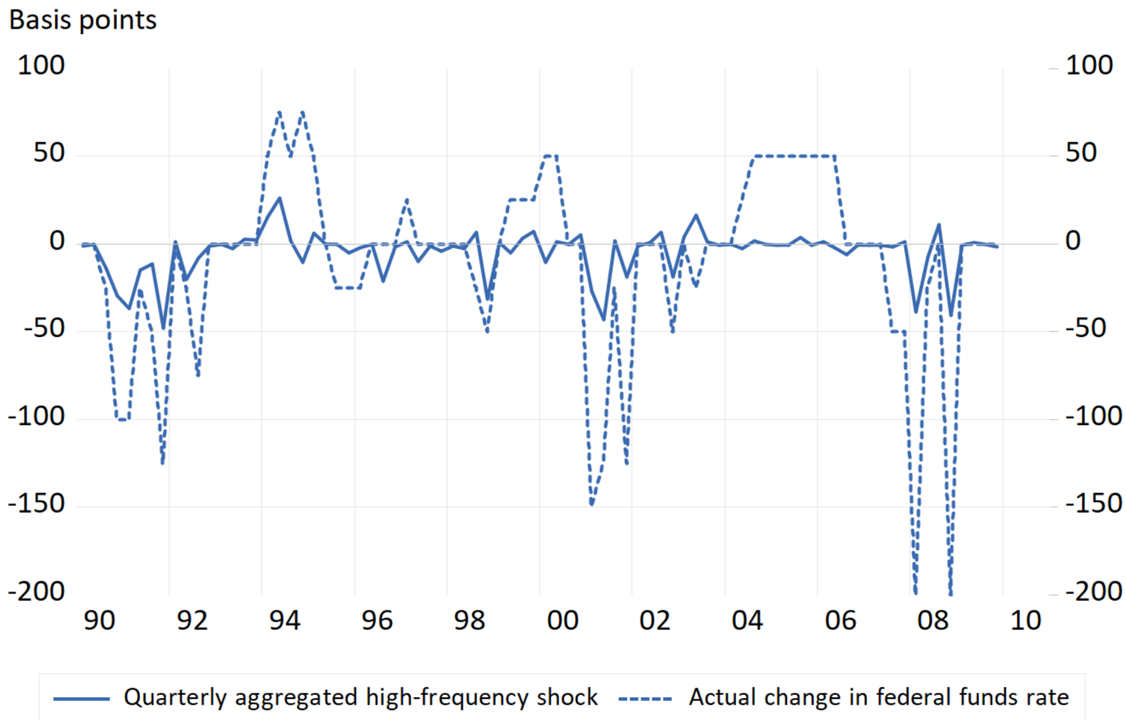

age group is likely a reduction in another age group. For both of these reasons, it is of interest to include all of the interaction terms in a single regression. As it is not practical to run a single regression that includes all of the age bins as regressors, we instead aggregate the age bins into three broad categories: 20-39, 40-59, and 60+. In this we are motivated by our finding that there are in essence three important age groups.

The results appear in Figure 12. The pattern is the same as in our baseline specification: A greater proportion of the young dampens the effects of a monetary policy shock and a greater proportion of the middle aged amplifies these effects. There is also weak evidence that a greater proportion of the old dampens the effects of monetary policy.

C.2 Aggregation to BEA Regions and Elimination of Small States

We test the robustness of our results to borders and the level of aggregation. One of the main differences between states and nations is that the borders of states tend to be more porous. It is therefore of interest to see if reducing the importance of borders affects our results. We run two

Figure 12: Wide Age Bins, Estimated Simultaneously

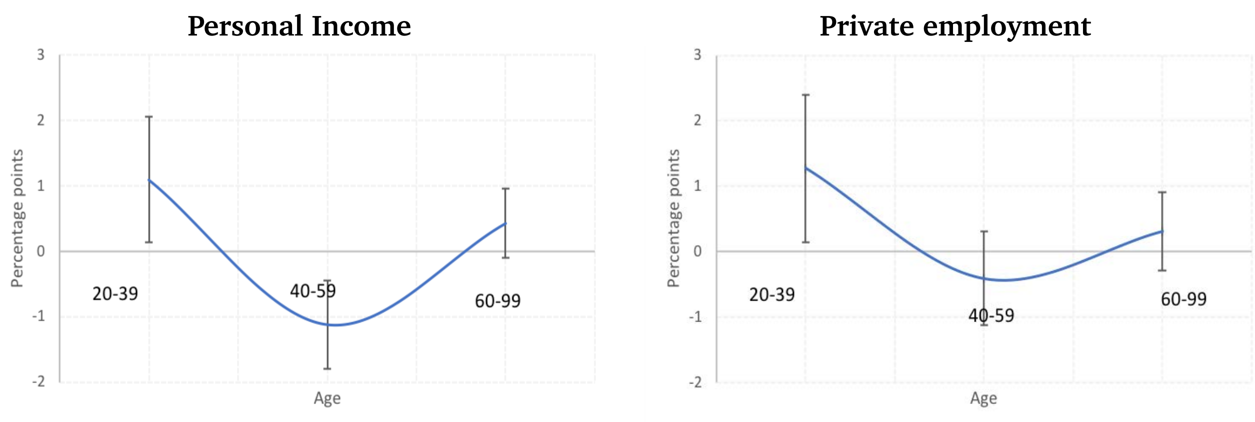

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

experiments. First, we aggregate states up to the level of BEA regions. This reduces the flow of people and goods into and out of the region as much of what was interregional trade and migration becomes intraregional. Second, we eliminate states with small populations, namely Alaska, North Dakota, Vermont, Washington D.C., and Wyoming. Migration may have a larger effect on these states.

Figure 13 presents the results. Aggregating to the 8 BEA regions reduces the degrees of freedom and hence the statistical significance of our results, but the basic pattern of young, middle aged and old remains. Dropping the small states has little effect on either the economic or statistical significance of our results.

C.3 Subsample stability

Federal funds futures markets began operating only in the late-1980's, which limits the extent of subsample stability that we can conduct using the high-frequency shocks. In addition, quarterly employment data is available at the state level beginning only in 1990. We compare our baseline results to two subsamples. First, we truncate all data in 2006 to avoid any contamination with the housing market crash and the Great Recession. Second, we begin our sample in 1995, after the FOMC began issuing press releases following every FOMC meeting.

Figure 13: BEA regions and Eliminate Small States

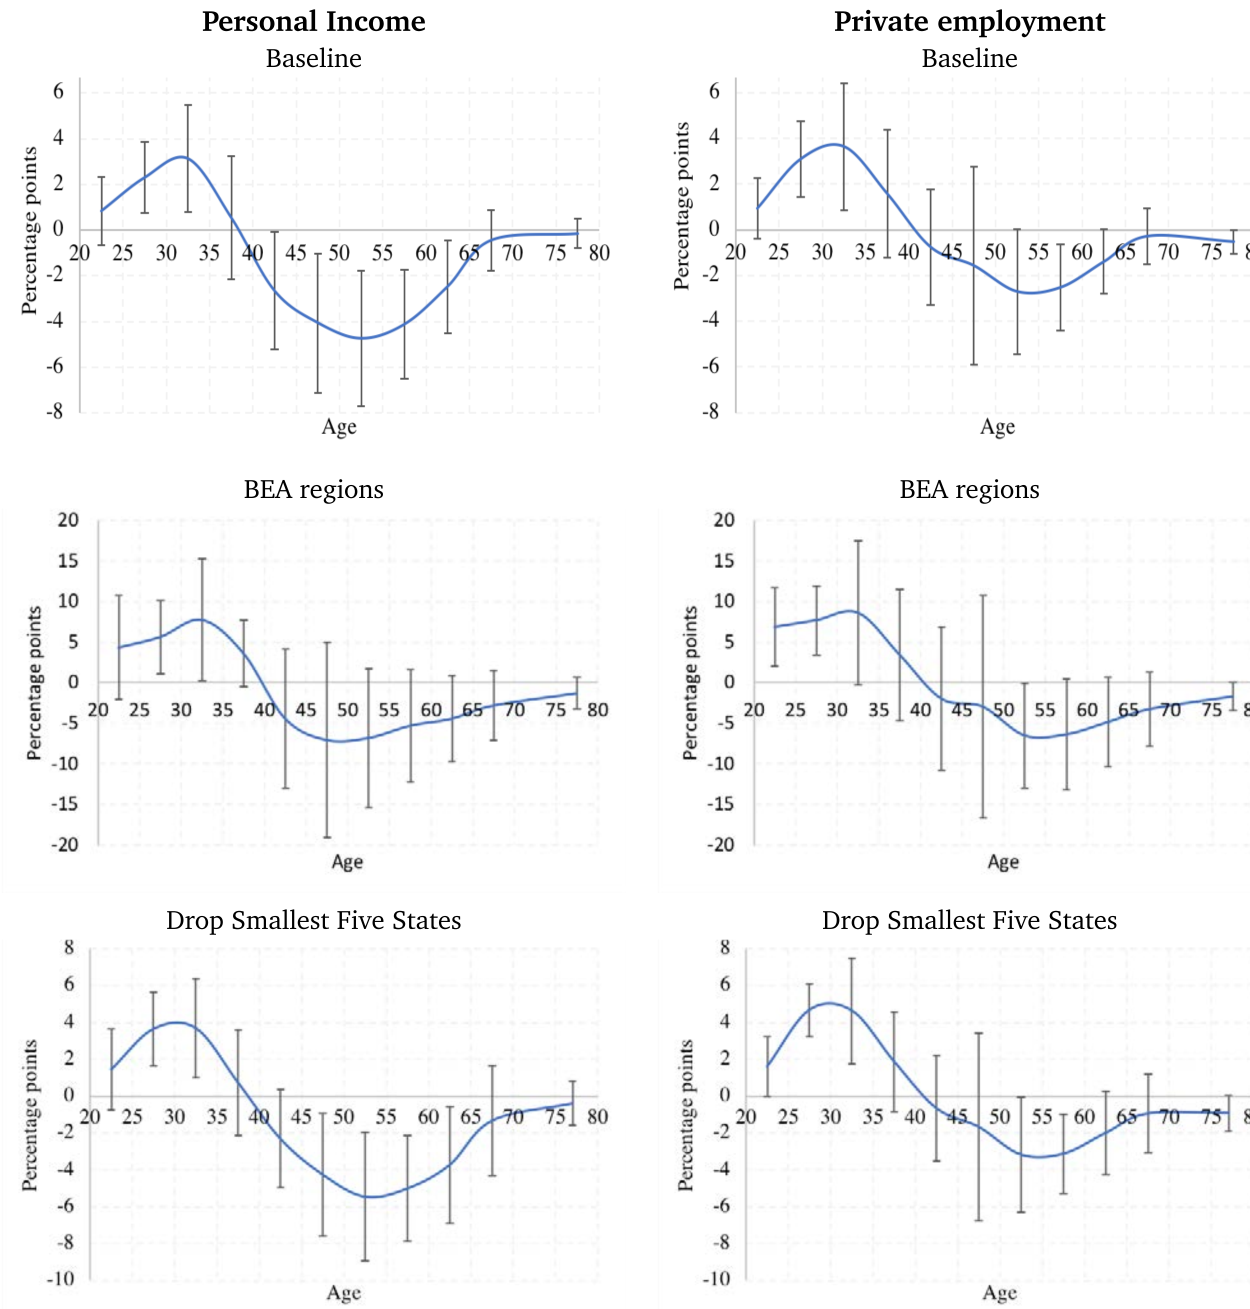

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 14 illustrates the effects on personal income and private employment in each of the subsamples. For ease of comparison, the first row contains our baseline results for the 1990-2008 sample period. The main story is largely unchanged for the different sample periods. Truncating the data in 2006 implies slightly weaker effects on both personal income and private employment than if we include the later period.

Beginning the estimation in 1995 exaggerates (amplifies) the demographic cycle for personal income. The dampening effect of the young and the stronger effect of the middle ages are both amplified in the 1995-2008 sample period as compared to our baseline results. For private employment, the effects from the younger age groups is smaller, but the effects stemming from the middle groups is even larger.

D State Characteristics

It is possible that age is simply picking up the effect of some other state characteristics. To control for state characteristics, we extend equation (1) by interacting control variables with the monetary policy shock,

$$\Delta log[X_{s,t+i}] = \alpha^{i,a}\epsilon_t^m p_{s,t}^a + \phi^{a,0}p_{s,t}^a + \beta^1 Z_s + \beta^2\epsilon_t^m Z_s + \gamma_s + \delta_t + u_{s,t}. \tag{5}$$

where Z_s refers to control variables, such as the fraction of manufacturing or construction employment, or percentage of employees with college degrees. We consider one additional control variable at a time.

Our first experiment is to take Z_s to be the state fixed effect, this will control for level differences in population across states that might be correlated with fixed state characteristics. [18] For example, Southern states may have more old and young but less industry. Figure 15, shows the results for this specification. The basic pattern is the same, although there is weak evidence that a greater fraction of older people may reduce the responsiveness of employment.

Interacting the shock with state fixed effects controls for a correlation between average population in the state and average characteristics of the state. The possibility remains that some

[18] Since we include state fixed effects (γ_s) in the regression, we only include Z_s interacted with the shock in this case.

Figure 14: Subsample Stability

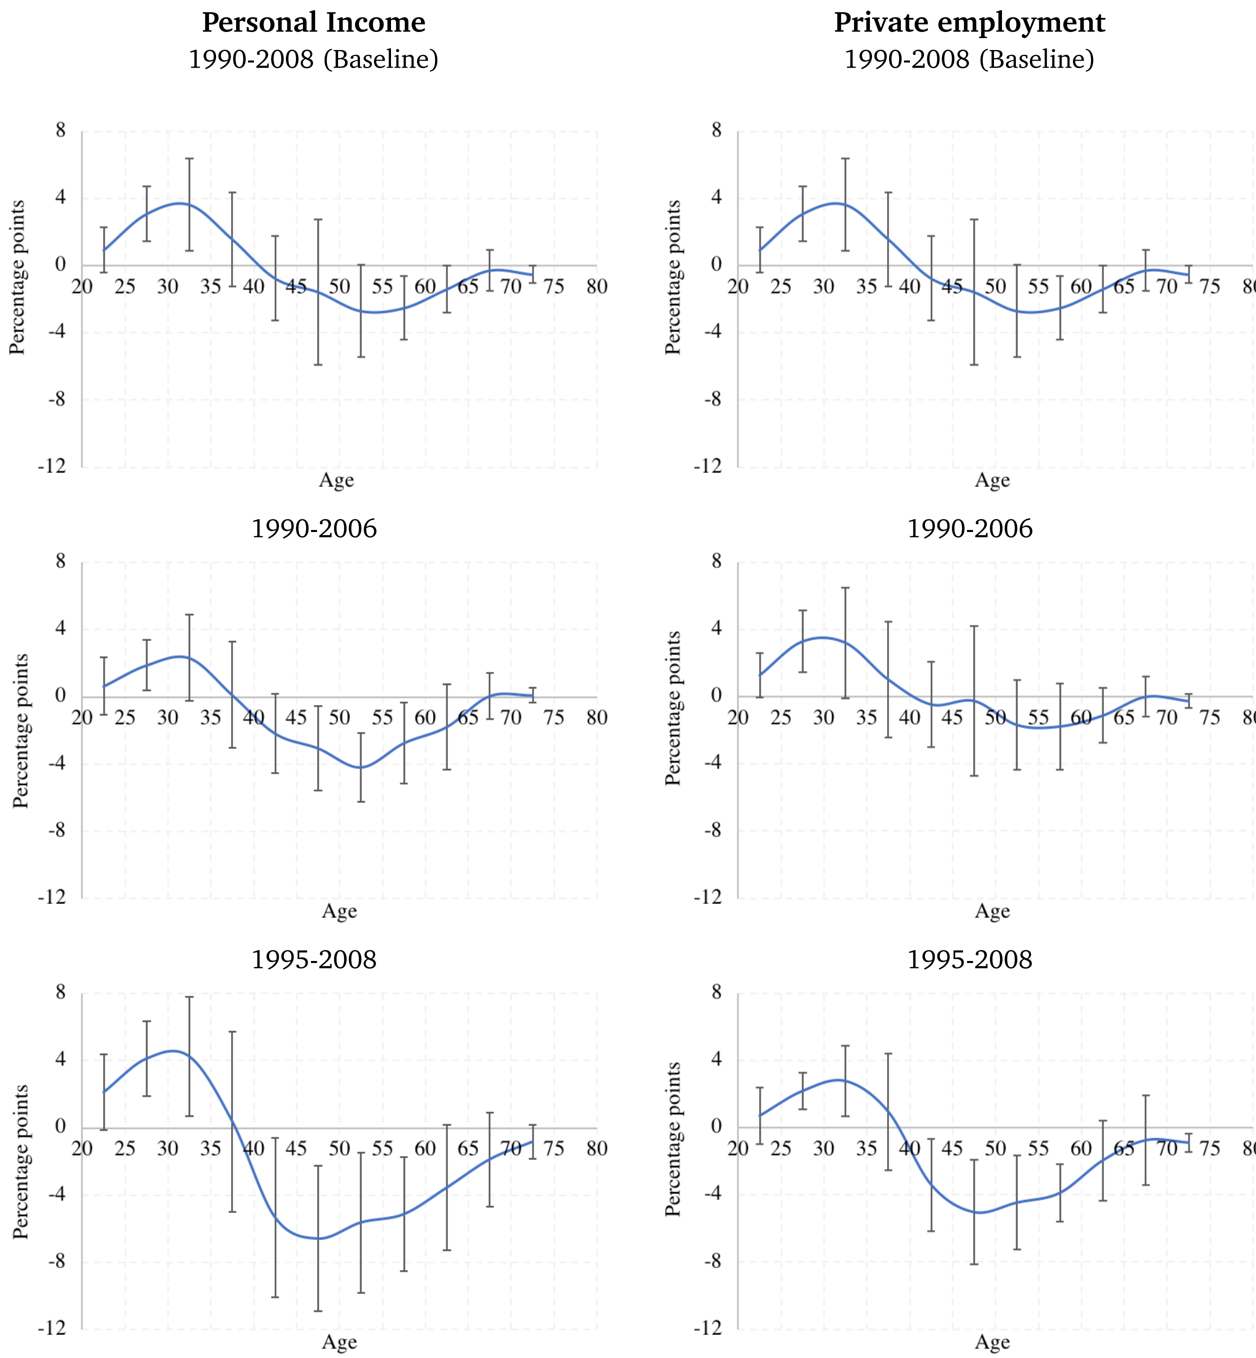

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 15: Include Shock Interacted With State

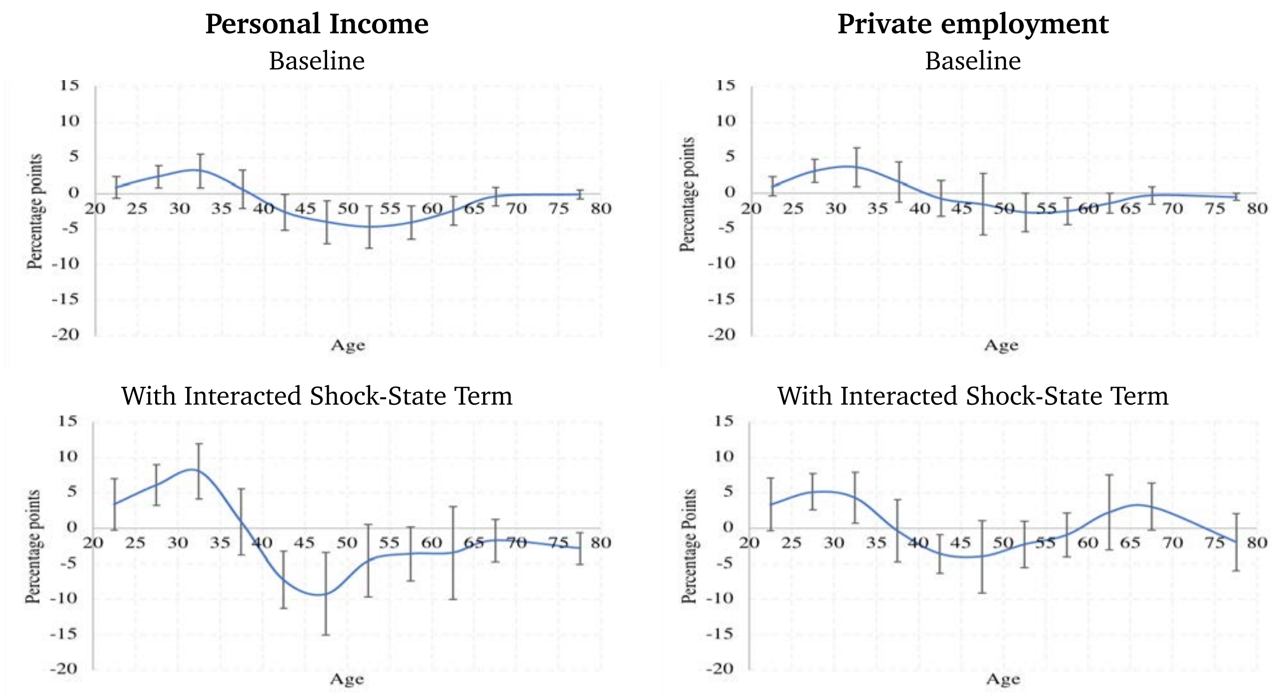

The horizontal axis represents the age groups. For a one-percentage point unexpected increase in the federal funds rate, the blue line represents the point estimates of the effect of a one-percentage point increase in the population share of group a on the growth rate of the variable of interest. Positive point estimates imply that, relative to the aggregate effects, the effects of monetary policy are weaker. Negative estimates imply stronger effects of monetary policy, relative to the aggregate effect. The vertical lines represent the 95% confidence intervals for the estimated coefficients.

characteristics may change over time in ways correlated with the change in demographic structure. We investigate a few possibilities. Figure 16 shows the results of including the fraction of sector employment to total private employment for manufacturing, finance, and construction in Z_s^1 and Z_s^2, that is both the fraction of employment and the fraction interacted with the monetary policy shock. Figure 16 presents the results. These additional regressors have little impact on the coefficients $\alpha^{i,a}$. In the case of construction, the results become stronger.

Next, we control for the income of the state by setting Z_s to be state personal income in equation (5). Interacting the shock with state personal income controls for differences in average income levels across states that might be correlated with the population structure. For example, incomes are lower for the young and old and higher for the middle aged. Figure 17 shows that controlling for the differences in incomes across states has a negligible impact on our results.

What else might age be proxying for? To separate out the effects of age from other demographic characteristics of states, we consider the percentage of employees who have college degrees, or are white, or are male. Interacting these characteristics with the monetary shocks allows us to separate out the responsiveness of these groups in the population to monetary policy. For example, if college educated workers are more informed about changes in and the effects of monetary policy, these workers might respond more to shocks relative to workers without college degrees. After controlling for this effect, the effects of the age distribution might disappear. Figure 18 presents these results. The results are largely unchanged. We conclude that obtaining a college degree, gender, and race are not driving our results.

Wong (2018) finds that the consumption of young homeowners respond more strongly to monetary policy shocks. To understand whether the housing market is driving our results, we control for house prices and home ownership rates. The results of this estimation are shown in Figure 19. The basic pattern survives.

E Firm size and age

Figures 20 and 21 illustrate the effects on income and employment broken down by firm size and age. Firm size in Figure 20 is measured by the number of employees. Firms of all types follow the

Figure 16: Control for Share of Employment in Other Sectors

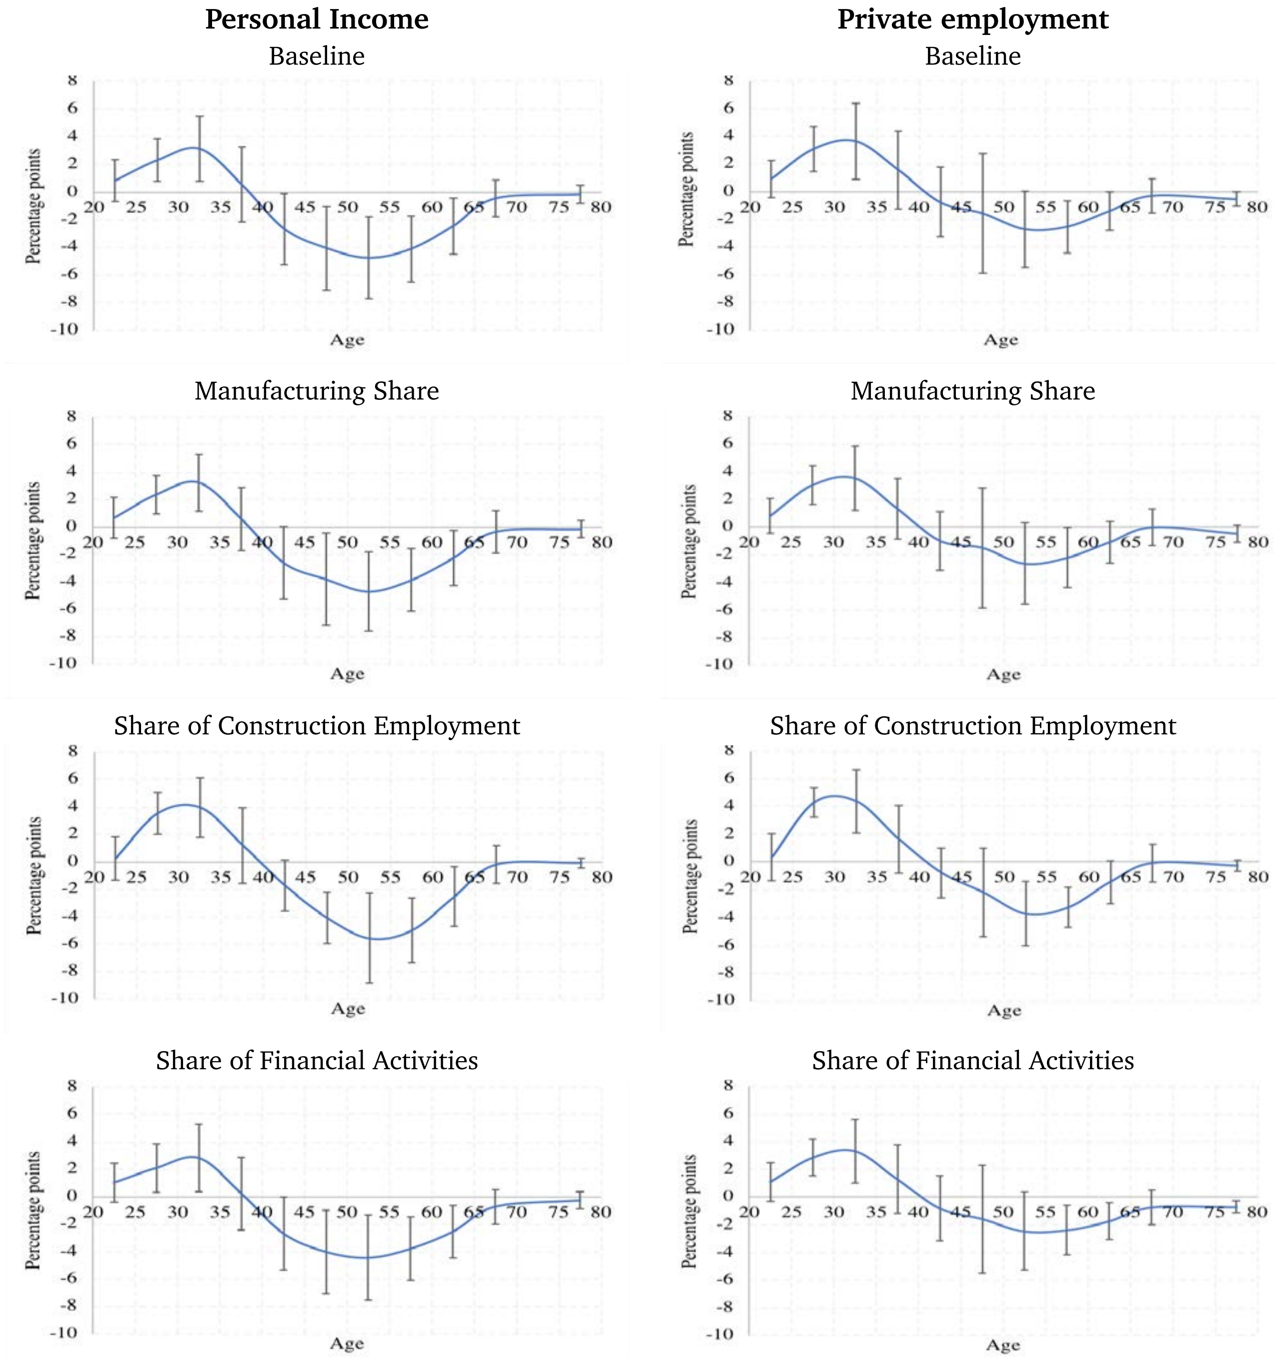

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 17: Control for State Personal Income (interacted with the Shock)

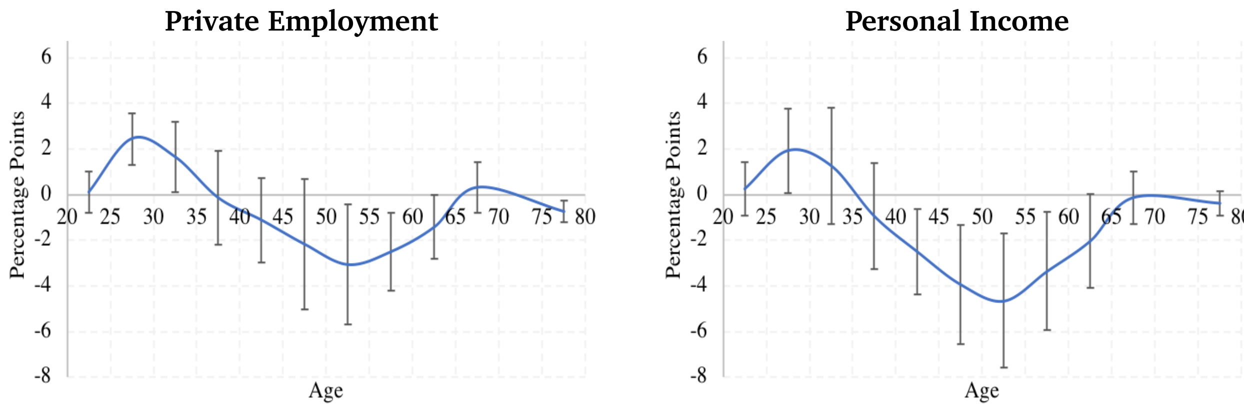

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

pattern in which the impact of monetary policy is dampened the greater the share of the young and amplified the greater the share of the middle aged. The effects are similar for employment, income per worker and total payroll, although the effects on total payroll are stronger than the effects on the number of jobs or earnings per worker. There is also some weak evidence that the pattern is stronger for the extensive margin (the number of jobs) than the intensive margin (earnings per worker) for young and small firms. There is weak evidence that the effects are exaggerated for younger and smaller firms.

Figure 22 depicts the results after we control for the share of small business. We include the fraction of firms with 249 or fewer employees interacted with the shock in the estimation. The pattern seen throughout the paper survives, indicating that correlation between the demographic structure and small businesses is not driving our results. The share of middle aged does not appear to be proxying for the share of small businesses in total employment.

Figure 18: Control for Other Demographic Characteristics

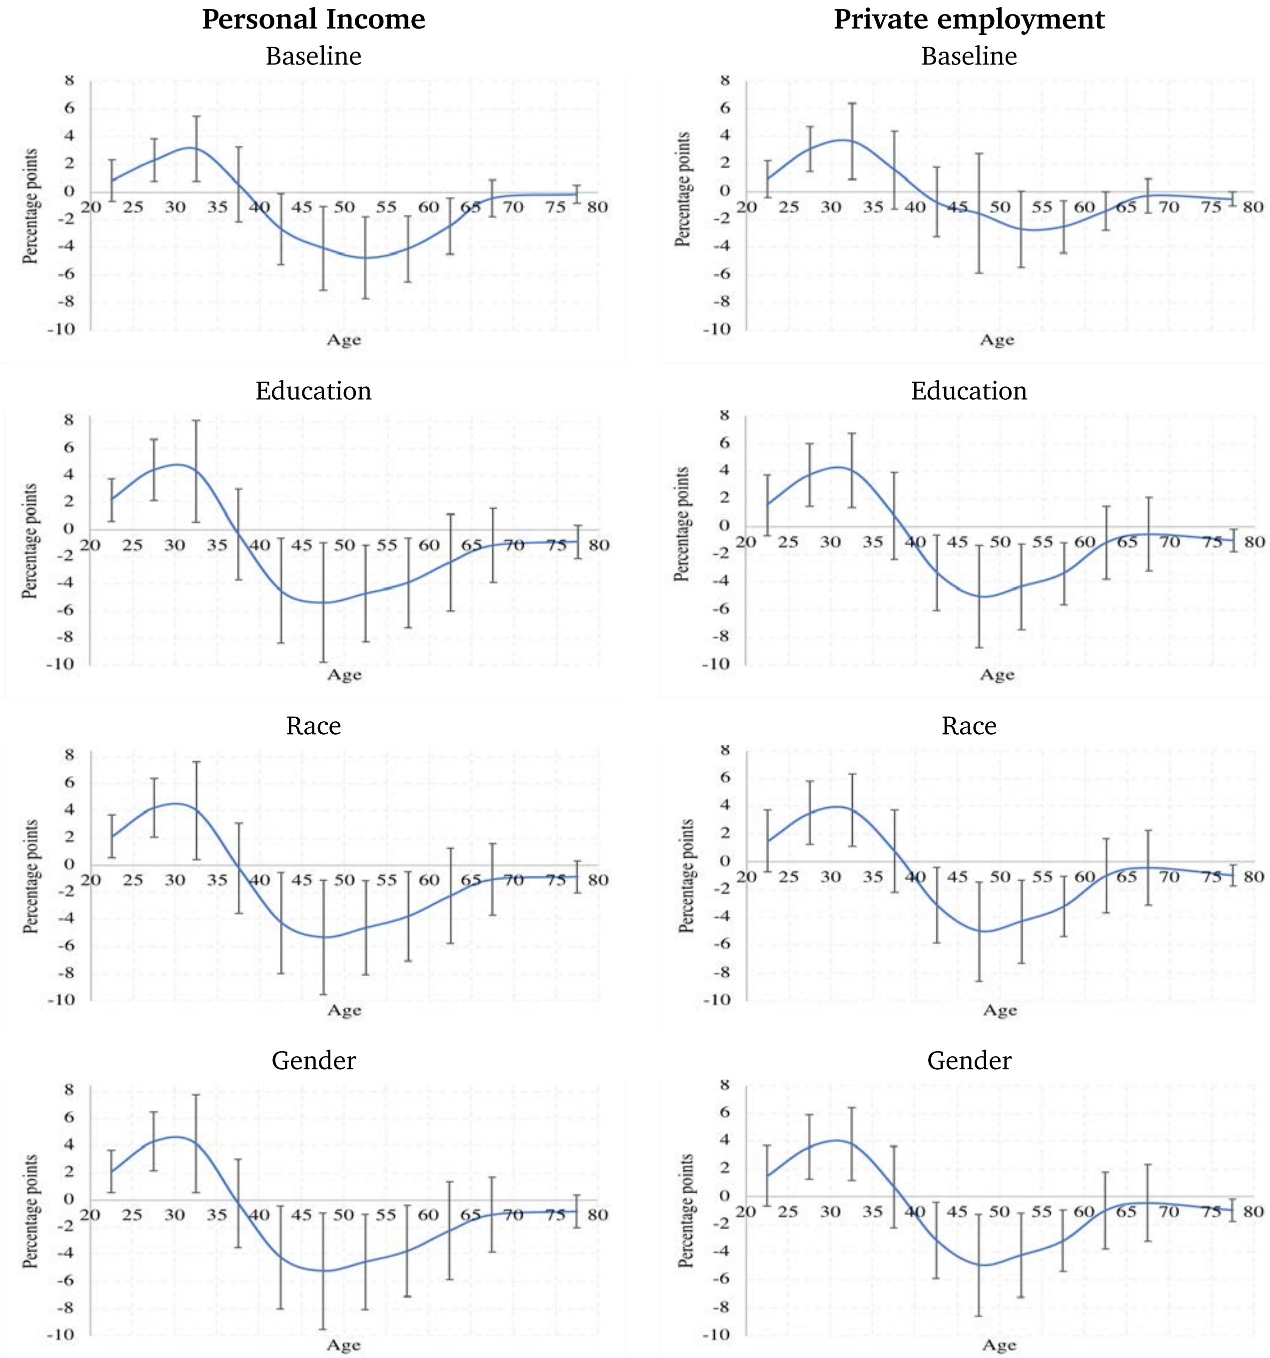

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 19: Control for Housing Related State-Level Characteristics

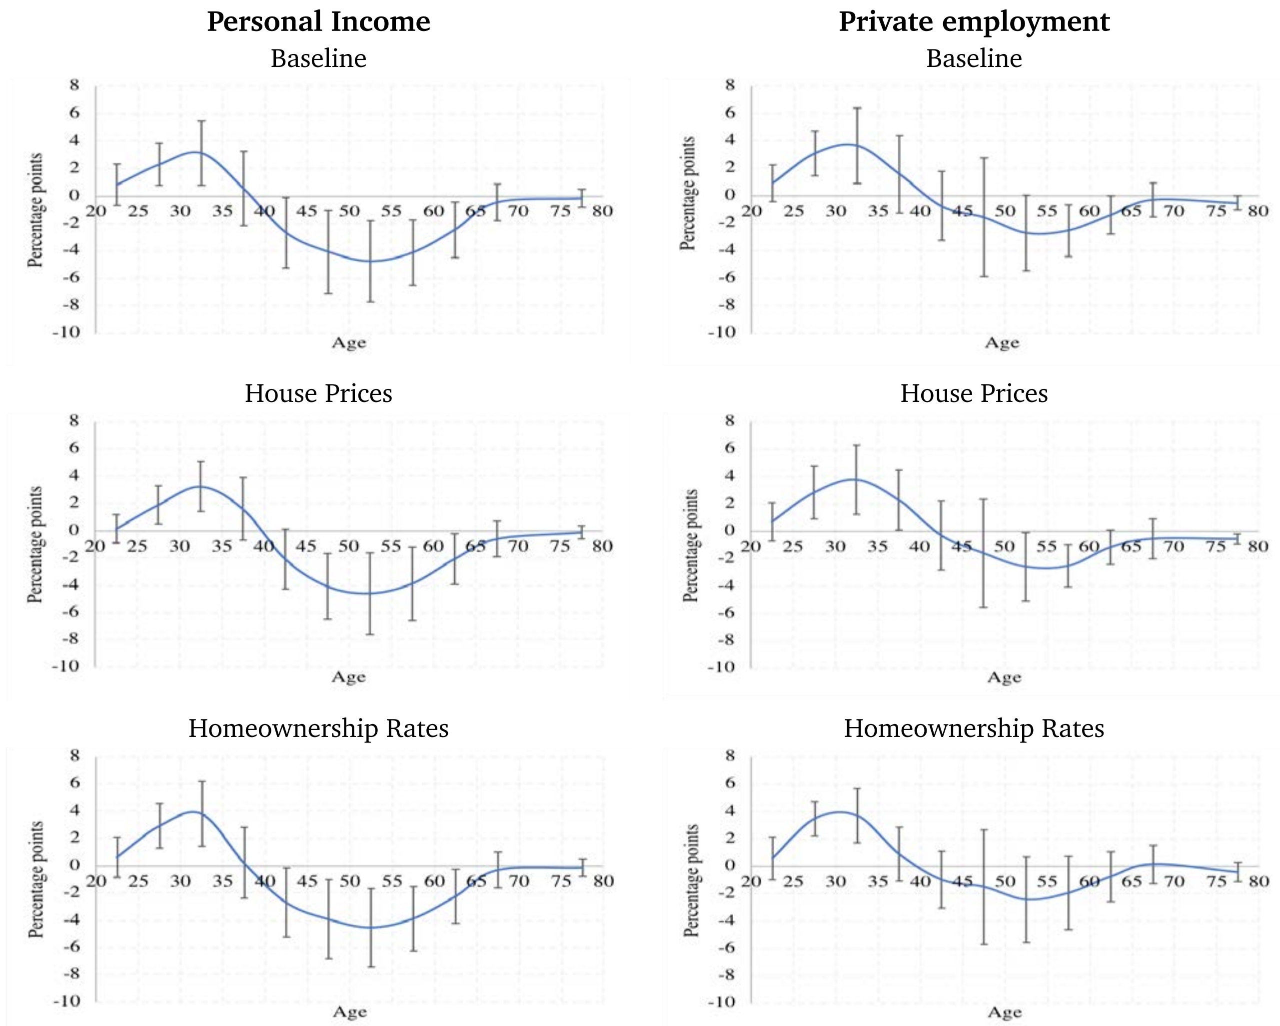

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 20: Response of Different Sized Firms

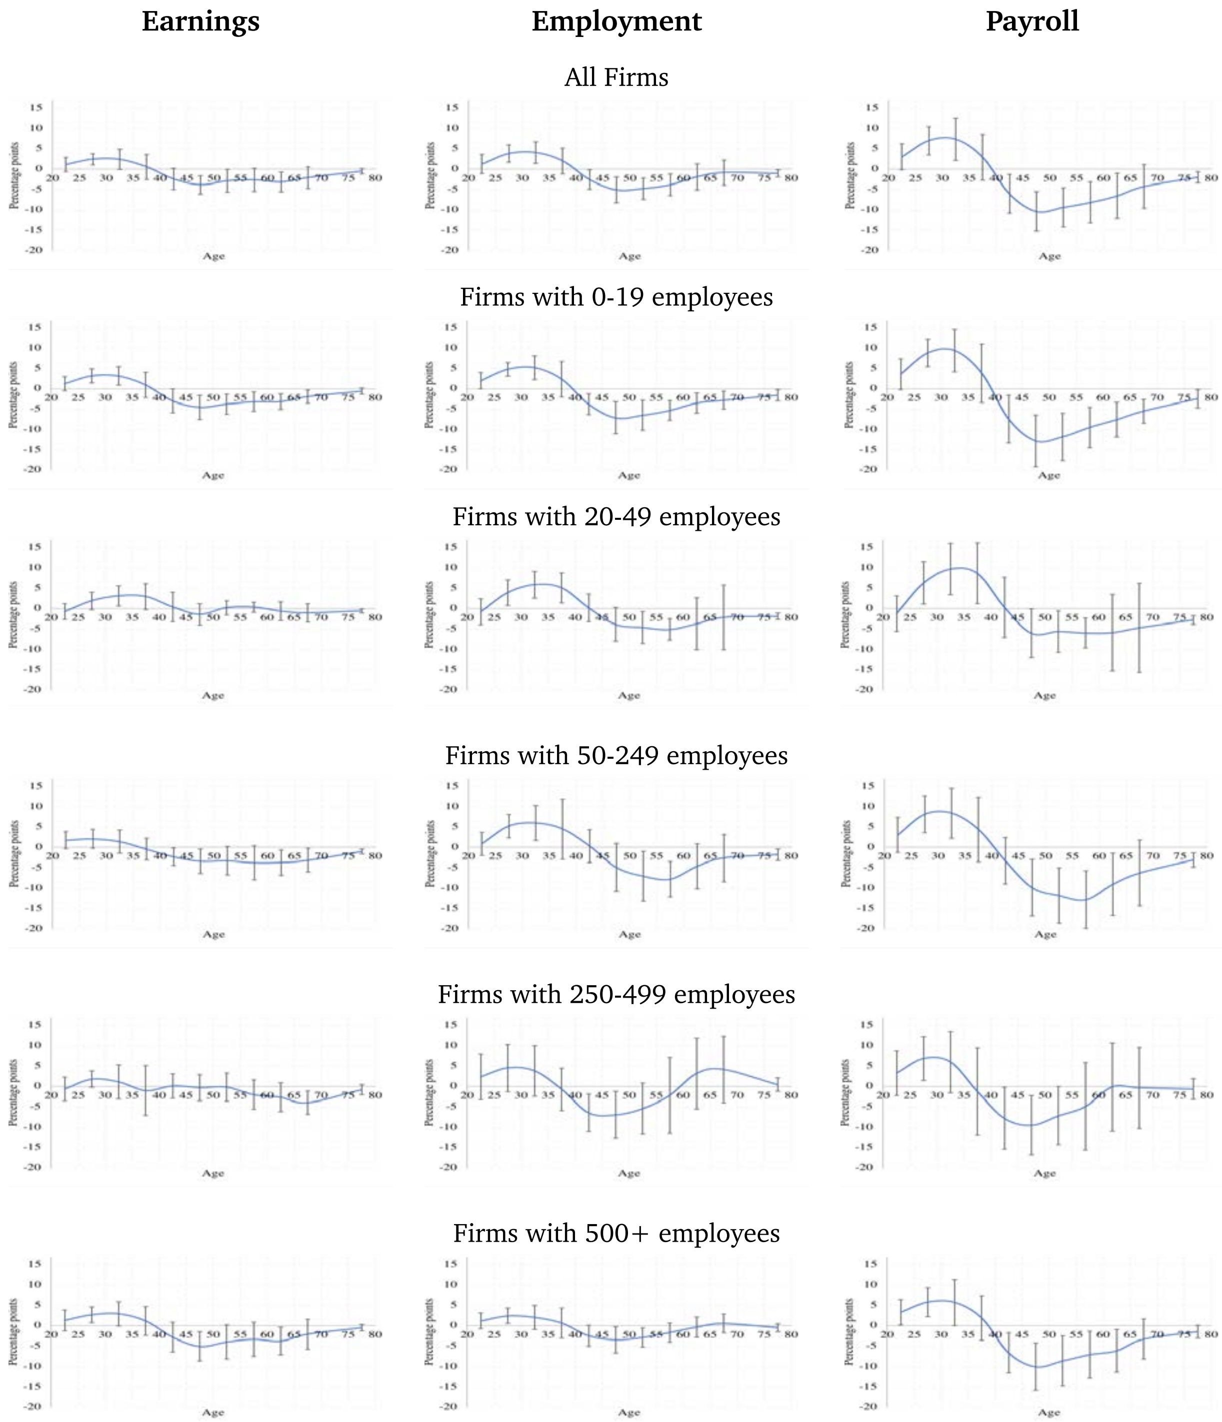

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 21: Response of Firms by Age

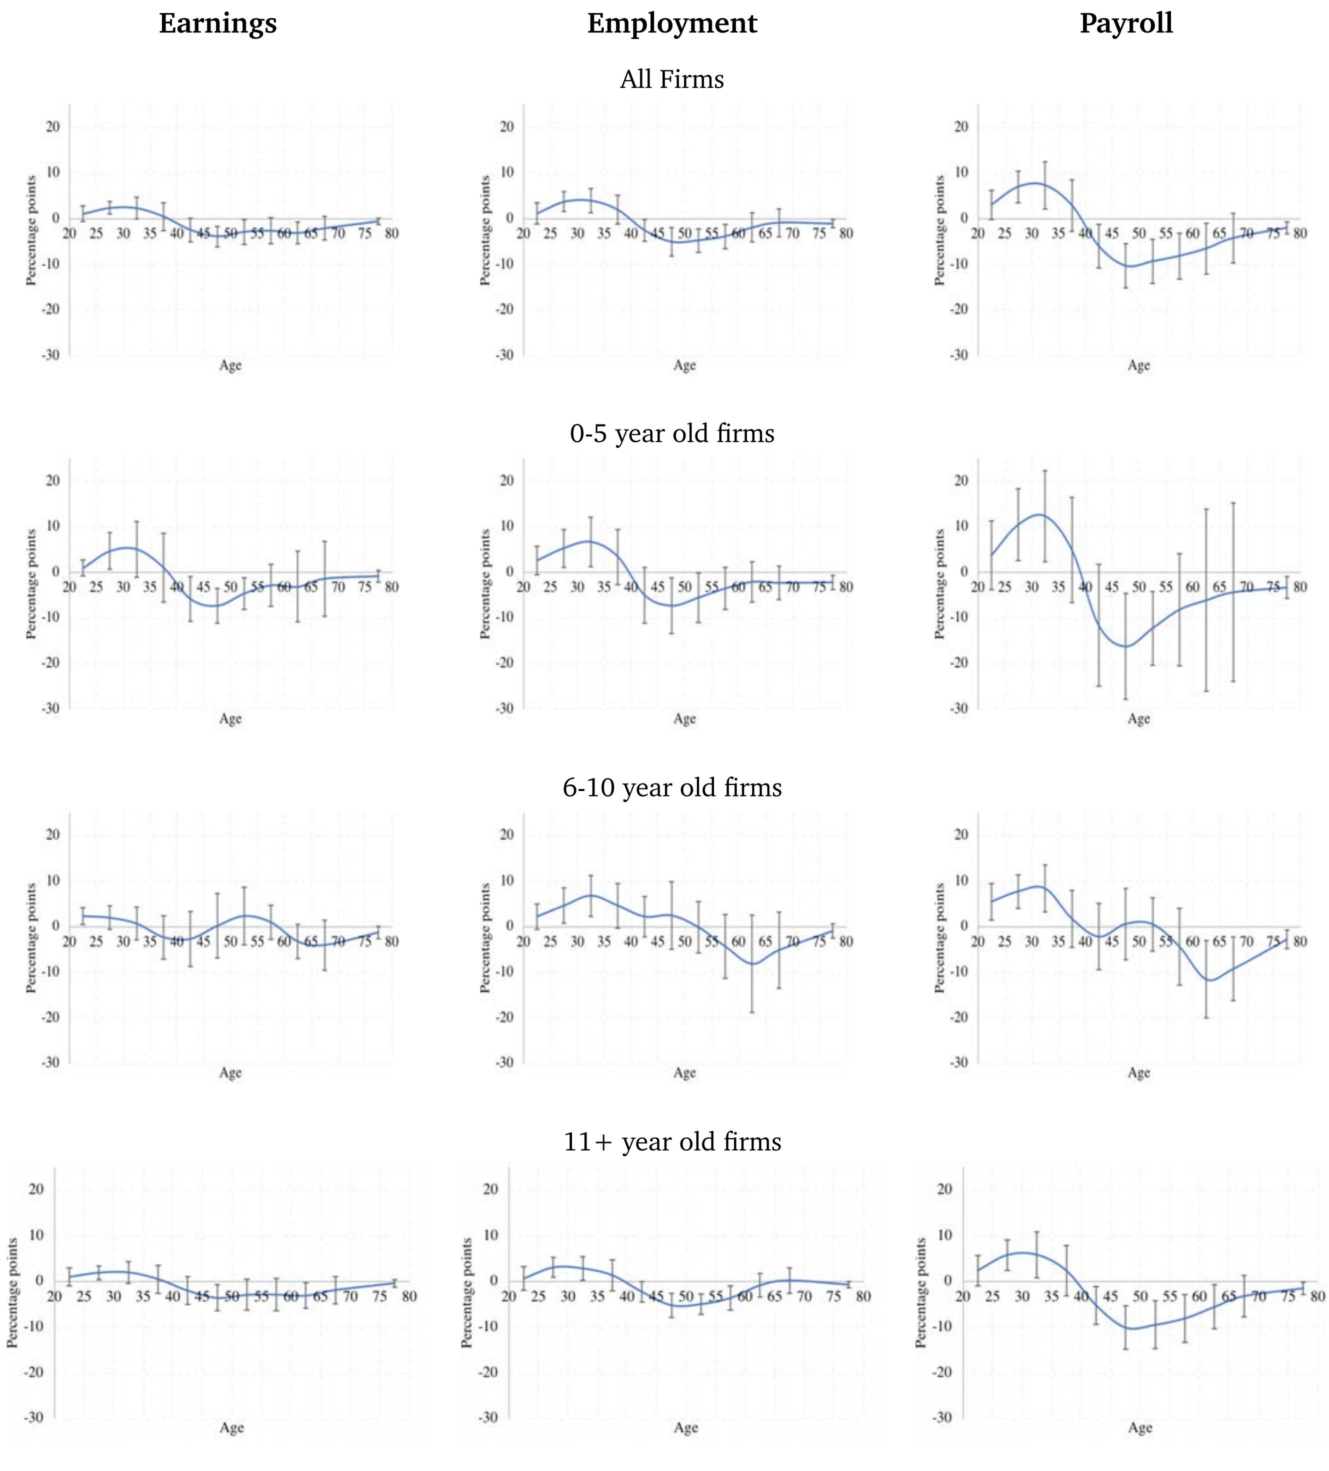

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.

Figure 22: Controlling for the Effect of Small Firms

Response of Personal Income (interact shock with fraction of small firms)

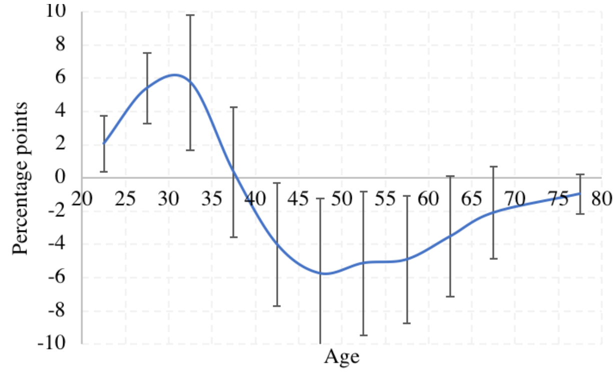

Response of Private employment (interact shock with fraction of small firms)

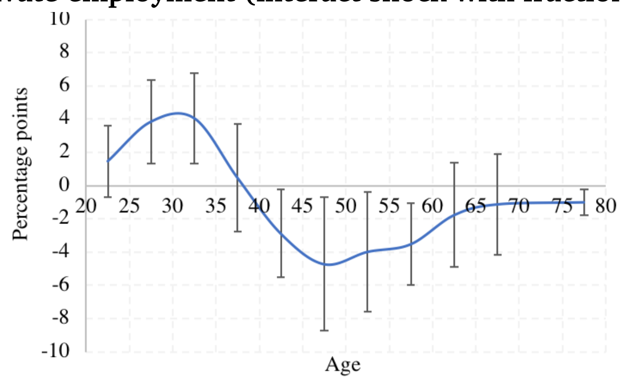

The horizontal axis represents the age groups a. The vertical axis represents the effect on the dependent variable of a positive one percentage point shock to the federal funds rate interacted with a one percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95% confidence intervals given Driscoll-Kraay standard errors.