American Economic Journal: Macroeconomics 2020, 12(2): 167–193
https://doi.org/10.1257/mac.20160419

The Aging of the Baby Boomers: Demographics and Propagation of Tax Shocks†

By Domenico Ferraro and Giuseppe Fiori*

We study how the changing demographic composition of the US labor force has affected the response of the unemployment rate to marginal tax rate shocks. Using narratively identified tax changes as proxies for structural shocks, we establish that the responsiveness of the unemployment rates to tax changes varies significantly across age groups: the unemployment rate response of the young is nearly twice as large as that of the old. This heterogeneity is the channel through which shifts in the age composition of the labor force impact the response of the unemployment rate to tax cuts. We find that the aging of the baby boomers considerably reduces the effects of tax cuts on aggregate unemployment. (*JEL* E24, E62, H24, H31, J21)

The post-World War II baby boom and the subsequent aging of the baby boomers resulted in dramatic shifts in the age composition of the labor force in the United States. Recent work has suggested that population aging is an important determinant of business cycle volatility (Jaimovich and Siu 2009) and of the propagation of monetary policy (Wong 2015). In this paper, we study how the changing age composition of the US labor force has affected the response of the unemployment rate to marginal tax rate shocks.

We proceed in two steps. First, we document that the unemployment rate response to tax shocks varies significantly across age groups. This heterogeneity is the channel through which shifts in the age composition of the labor force affect the response of the unemployment rate to tax changes. Second, we quantify the extent to which the aging of the baby boomers has changed the effects of marginal tax rate shocks on the aggregate unemployment rate.

* Ferraro: Department of Economics, W.P. Carey School of Business, Arizona State University, PO Box 879801, Tempe, AZ 85287 (email: domenico.ferraro@asu.edu); Fiori: Department of Economics, Poole College of Management, North Carolina State University, Box 8110, Raleigh, NC 27695 (email: gfiori@ncsu.edu). Richard Rogerson was coeditor for this article. We thank an anonymous referee, Alexander Bick, Carlo Favero, John Haltiwanger, Berthold Herrendorf, Bart Hobijn, Nir Jaimovich, Benjamin Moll, Pietro Peretto, Roberto Pinheiro, Ed Prescott, Valerie Ramey, Todd Schoellman, Matthew Shapiro, Pedro Silos, Nora Traum, Etienne Wasmer, seminar participants at Federal Reserve Board, USC, UC Riverside, Annual Meeting of the Society for Economic Dynamics, Annual Congress of the European Economic Association, Macro Workshop at Arizona State University, Triangle Dynamic Macro Workshop at Duke University, and OFCE-SKEMA Workshop on Economic Growth, Innovation, and Corporate Governance for comments and discussions.

† Go to https://doi.org/10.1257/mac.20160419 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

To estimate the responses of age-specific unemployment rates to tax changes, we rely on narrative identification of tax shocks (Romer and Romer 2009, 2010). Specifically, we use narratively identified tax changes as proxies for structural tax shocks, and structural vector autoregressions (proxy SVARs) to estimate the dynamic responses to a tax cut (Mertens and Ravn 2013, Mertens and Montiel Olea 2018). Using CPS data, we construct average marginal tax rates (AMTRs) by age groups, as we are interested in the transmission mechanism of tax changes that operate via incentive effects on intertemporal substitution. In addition, we build a new set of age-specific proxies to account for the fact that tax reforms may have changed AMTRs of different age groups differently. Indeed, our measurement points to substantial age heterogeneity in AMTRs and how tax reforms impacted AMTRs of different age groups. Proxy SVARs are estimated separately for the young (16–34 years old), prime-age (35–54 years old), and old (55 years and older), using our constructed age-specific proxies as external instruments for age-specific AMTRs.

We establish that the unemployment rate response of the young is nearly twice as large as that of the old. By contrast, we show that the age-specific labor force shares are in fact unresponsive to the (narratively identified) changes in AMTRs. This empirical finding is, perhaps, not surprising as the observed demographic trends in the age composition of the labor force in the United States are largely determined by fertility decisions made long before a specific tax shock. Thus, the age composition of the labor force is largely predetermined at the time of a legislated tax change. In addition, the post-war baby boom and the aging of the baby boomers in the last thirty years resulted in large movements in the age composition of the labor force. As a result, time-series variation in labor force shares by age abounds.

We use these two pieces of evidence to measure how the responsiveness of the US unemployment rate to tax changes depends on the age composition of the labor force. To this aim, we construct an aggregate unemployment response to tax shocks, that accounts for the observed movements in the age composition of the labor force. The implied response provides a simple quantitative accounting of how the observed demographic trends in the United States impact the effectiveness of tax cuts in reducing unemployment. Specifically, when an economy is characterized by a smaller share of young workers, everything else equal, one observes a smaller aggregate response to tax cuts. We find that the aging of the baby boomers reduces the response of the unemployment rate to tax cuts by 40 percent.

The implications for fiscal policy are far-reaching. In the United States, given the current fertility and mortality rates, the working-age population is expected to become older. Similar estimates and projections apply to Japan and most industrialized countries in Europe. The results in this paper indicate that tax shocks of the size observed in the United States since World War II are becoming increasingly less effective in stimulating economic activity.

Related Literature.—The debate in the aftermath of the Great Recession of 2007–2009 has led to renewed interest in the question of how fiscal policy affects the economy. As a result, a growing strand of the empirical literature investigates the effects of government purchases and taxes (Romer and Romer 2010, Barro and

Redlick 2011, Ramey 2011, Mertens and Ravn 2013, Mertens and Montiel Olea 2018).[1] This literature considers aggregate macroeconomic variables, which is the natural starting point for analyzing the economic forces that shape the aggregate response to fiscal shocks. Here, we pursue a disaggregated analysis by considering one specific dimension of heterogeneity: age. This approach sheds light on the link between microeconomic behavior and macroeconomic effects of tax changes. We emphasize, however, that our ultimate goal is to gauge the implications of demographic change for the aggregate labor-market response to tax changes. To date, this paper is the first attempt to tackle this question.[2]

Recent work has studied the implications of demographic change for aggregate labor-market dynamics. For instance, Shimer (1999) shows that the entry of the baby boomers into the labor force in the late 1970s, and their aging, accounts for most of the low-frequency movements in the US unemployment rate since World War II. Jaimovich and Siu (2009) show that demographic change accounts for a significant fraction of the decrease in business cycle volatility observed in the United States since the mid-1980s. In this paper we argue that the aging of the baby boomers considerably reduces the effects of tax changes on aggregate unemployment.

Furthermore, we argue that assessing the effects of tax changes across different age groups is also relevant for distinguishing between competing transmission mechanisms of tax shocks. Understanding if and to what extent the young, prime-age, and old display differences in the unemployment responsiveness to tax shocks would seem important for understanding why aggregate unemployment responds to tax changes as much as it does. Analogously, Ríos-Rull (1996); Gomme et al. (2005); Hansen and İmrohoroğlu (2009); and Jaimovich, Pruitt, and Siu (2013) assess the implications of age-specific differences in cyclical movements of hours worked.

I. Basic Facts on Population Aging

In this section we provide a bird's-eye view of the aging of the workforce observed in the United States in the last thirty years. The post-WWII baby boom and the subsequent baby bust resulted in dramatic shifts in the age composition of the working-age population, labor force, employment, and unemployment. To the extent that the young, prime-age, and old feature different labor force attachment, turnover rates, and job search intensities, population aging has important implications for the aggregate labor market response to tax changes. These observations motivate our paper.

Working-age Population.—Panels A and B in Figure 1 report the average age and the shares by age of the US working-age population, respectively. Two patterns

[1] See Ramey and Shapiro (1998); House and Shapiro (2006, 2008); Pappa (2009); Brückner and Pappa (2012); Auerbach and Gorodnichenko (2012); Favero and Giavazzi (2012); Cloyne (2013); Ramey and Zubairy (2018); Nakamura and Steinsson (2014); Acconcia, Corsetti, and Simonelli (2014); Mertens and Ravn (2014); and Caldara and Kamps (2017) for further references.

[2] Anderson, Inoue, and Rossi (2016) document heterogeneous effects of government spending shocks on consumption, depending on income and age.

Figure 1. Trends in the Age Composition of the US Population, 1950–2015

Notes: Panel A shows the average age of the US civilian noninstitutional population (20–64 years old). The average age of the population is calculated as $\bar{a}^P \equiv \sum_{a \in A} \left((\underline{a} + \bar{a})/2 \right) \phi_a^P$, where $\underline{a}$ and $\bar{a}$ are respectively lower and upper bounds of the age group $a \in A$, with $A = \{20\text{–}24, 25\text{–}34, 35\text{–}44, 45\text{–}54, 55\text{–}64\}$, and ϕ_a^P is the age-specific population share (the ratio of the population in the age group a to total population). Panel B shows the population shares by three age groups: (i) full line with circles (left axis) shows $\phi_{20\text{–}24}^P + \phi_{25\text{–}34}^P$; (ii) dashed line with squares (left axis) shows $\phi_{35\text{–}44}^P + \phi_{45\text{–}54}^P$; and (iii) dashed-dotted line with diamonds (right axis) shows $\phi_{55\text{–}64}^P$.

emerge. First, population has been aging since the mid-1980s. Second, the average age of the population has been declining over the course of twenty-five years from the early 1960s to the mid-1980s as a result of the sharp increase in birth rates after World War II. This is the so-called “baby boom.” In the early 1960s, birth rates started to decline toward levels prior to the baby boom, which led to the subsequent “baby bust.” Since the mid-1980s, the population has been aging, fueled by the aging of the baby boomers.

These slow-moving trends resulted in large shifts in the age composition of the population. The share of the 20–34-year-olds increased by 10 percentage points (from approximately 35 to 45 percent) over twenty years from 1960 to 1980. By contrast, during the same period, the population share of the 35–54-year-olds declined by nearly the same amount. The population share of the 55–64-year-olds remained approximately constant over the same period. It starts declining in the early 1980s to reach a trough in the mid-1990s. At that time, it sharply reverts to steady growth, reaching roughly 22 percent of the overall population in 2015. The share of the 20–34-year-olds in the working-age population starts its decline in the mid-1980s, reaching a plateau in the early 2000s.

Labor Force (Employment Plus Unemployment).—The demographic trends observed in the working-age population have led to changes in the age composition and thus the average age of the labor force. Notably, the labor force has been aging since the mid-1980s, as has the working-age population (see Figure OA.1 in the online Appendix). The share of the 20–34-year-olds in the labor force increased by more than 10 percentage points (from approximately 35 to 48 percent) over fifteen years from the mid-1960s to early 1980s. During the same period, the share of the 35–54-year-olds declined by approximately 10 percentage points, whereas the labor force share of the 55–64-year-olds started declining in the early 1970s, reached a trough of 10 percent in the mid-1990s, and has steadily increased since then.

The average age of employed individuals tracks closely with that of the labor force (see Figure OA.2 in the online Appendix). This observation is not surprising as employment represents nearly 95 percent of the labor force for 1950–2015. Movements in the average age of the unemployed, and shifts in their unemployment shares, are substantially larger than those in employment (see Figure OA.3 in the online Appendix). In the mid-1950s, the average age of an unemployed worker was 39 years, whereas in the 1980s, the average age was as low as 33. Yet, as of 2015, an unemployed person is, on average, nearly the same age as an unemployed person in the years preceding the baby boom. The unemployment share of the 20–34-year-olds rose by nearly 20 percentage points (from approximately 45 to 65 percent) from the mid-1960s to the 1980s. Since 1980, the share of the 20–34-year-olds has been declining at a nearly constant pace. As of 2015, they represent 49 percent of the unemployment pool. During the same period, the unemployment share of the 35–54-year-olds has instead declined and then risen, mirroring the unemployment share of the 20–34-year-olds. Finally, the unemployment share of the 55–64-year-olds declined from the mid-1950s to the 1990s, and it has been steadily rising since then. As of 2015, the share of the 55–64-year-olds constitutes 14 percent of unemployed persons.

II. Marginal Tax Rates by Age

In this section we describe the methodology to construct average marginal tax rates (AMTRs) by age groups. Based on Barro and Redlick (2011), and most of the literature thereafter, we consider a notion of labor income that includes wages, self-employment, partnership, and S-corporation income. Data are taken from the CPS March Supplement that allows us to match individual income and demographic characteristics, such as age. AMTR is the sum of the federal individual income tax and the payroll (FICA) tax. We use NBER-TAXSIM to simulate marginal income tax rates and marginal payroll tax rates at the individual level. We then construct AMTRs by age groups as the sum of average marginal individual income tax rates (AMIITRs) and average marginal payroll tax rates (AMPTRs), using adjusted gross income shares as weights. As such, we compute a hitherto unexplored measure of AMTRs conditioning on the age composition of taxpayers. (We refer the reader to Appendix A for more details about the construction of the AMTRs.)

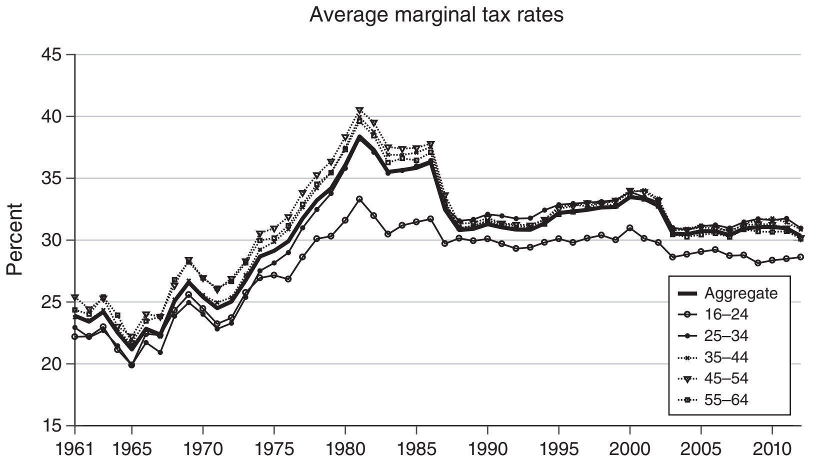

Figure 2. Average Marginal Tax Rates, 1961–2012

Notes: The figure shows the time series of average marginal tax rates (AMTRs), aggregate and by age groups. AMTR includes the average marginal individual income tax rate (AMIITR) and average marginal payroll tax rate (AMPTR).

Figure 2 reports aggregate and age-specific AMTRs for the sample period 1961–2012.[3] Three patterns emerge. First, aggregate AMTRs display a marked upward trend from the early 1960s to the early 1980s. It fluctuates in the 24–27 percent range over roughly ten years from 1961 to 1970. In the 1970s, AMTRs sharply rise from 25 percent toward the post-war peak of 38 percent in the early 1980s. This acceleration was primarily due to the bracket creep effects induced by rising inflation over the so-called Great Inflation of the 1970s. After the 1980s, the sustained rises in the FICA tax have been almost entirely offset by reductions in the federal individual income tax rates, which have remained in the 20–25 percent range since then. In addition to these long-run trends, the time series of AMTRs features substantial year-to-year variation with a standard deviation of 4.1 percent. As discussed in Mertens and Montiel Olea (2018), the bulk of this year-to-year variation is driven by statutory changes in federal individual income taxes. Consistent with the literature, AMTRs do not include state-level taxes. However, the amount of short-run variation in state-level marginal tax rates is small (see Barro and Redlick 2011).

Second, AMTRs differ by age, albeit less than one would have expected a priori. The exception is the age group 16–24 with AMTRs that are considerably below the aggregate; their average AMTR is 28 percent compared with the aggregate AMTR of 30.3 percent. The average AMTR is 30.1 percent for the group 25–34 and 30.8 percent for the groups 35–44 and 55–64. Third, the dynamics of the tax rates are remarkably similar across the different age groups, tracking closely the aggregate AMTR, with the lowest correlation being 0.93.

[3] Our sample is constrained by CPS data as its availability starts in 1961.

III. Econometric Methodology

In this section we discuss the identification of tax shocks, specification, and estimation.

A. *Identification*

We use structural vector autoregressions (SVARs) to measure the dynamic effects of changes in AMTRs. SVARs have been extensively used in macroeconomics to evaluate the effects of fiscal policy actions as well as other aggregate shocks, such as monetary policy, technology, and oil price shocks. (See Ramey 2016 for a survey of the literature.) As in the SVAR tradition, we associate a "tax shock" to a VAR innovation to the AMTR that jointly satisfies three criteria: (i) it is unpredictable, given current and past information; (ii) it is uncorrelated with other structural shocks; and (iii) it is unanticipated, that is, it is not news about future policy actions.

As shown in Figure 2, the post-WWII history of US federal income tax policy includes several large increases and decreases in marginal tax rates, which arguably provides valuable identifying variation. Yet, the vast majority of the observed legislated changes in tax rates result from policy actions aimed at offsetting cyclical downturns. This poses well-known challenges for the identification of the causal effects of tax changes on aggregate outcomes.

To address this issue, identification of tax shocks is obtained by SVARs and proxies for exogenous variation in tax rates as external instruments (see Mertens and Ravn 2013). To this goal, we construct time series of age-specific proxies for exogenous changes in age-specific AMTRs. Our strategy follows Mertens and Montiel Olea (2018). To select instances of exogenous variation in tax rates, they rely on the narrative approach of Romer and Romer (2009): changes in total tax liabilities are classified as "exogenous" based on the motivation for the legislative action being either long-run considerations that are unrelated to the business cycle or inherited budget deficits. An additional concern is that legislated tax changes are often implemented with a considerable lag, which can generate anticipation effects. Indeed, Mertens and Ravn (2012) provide evidence of aggregate effects of legislated tax changes prior to their implementation. To avoid these anticipation effects, we consider observations on individual income tax liability changes legislated and implemented within the year as in Mertens and Ravn (2013). Based on these considerations, Mertens and Montiel Olea (2018) identify seven exogenous tax reforms.

In Table 1, we report how each of these tax reforms changed the aggregate AMTR. Specifically, the impact of a reform is measured as the difference between two counterfactual tax rates. The first counterfactual tax rate is calculated using the year $t-1$ income distribution and year t statutory tax rates and brackets. The second counterfactual tax rate is calculated based on the year $t-1$ income distribution and year $t-1$ statutory tax rates and brackets. The difference between the two, then, isolates the impact that a tax reform implemented in year t had on the AMTR.[4] An

[4] The calculations shown in the column labeled "FF" are based on individual-level simulated tax returns from the NBER-TAXSIM program, whereas the figures from Mertens and Montiel Olea (2018) labeled "MMO" are

Table 1—Estimated Impact of Selected Tax Reforms

		Impact on AMTRs (percentage points)	
	In year	MMO (1)	FF (2)
Revenue Act of 1964	1964	−2.61	−2.26
Revenue Act of 1978	1979	−1.35	−1.91
Economic Recovery Tax Act 1981	1981	−0.31	−0.36
Tax Reform Act of 1986	1987	−2.41	−4.18
Omnibus Budget Reconciliation Act of 1990	1991	0.79	0.44
Omnibus Budget Reconciliation Act of 1993	1993	1.08	0.12
Jobs and Growth Tax Relief Reconciliation Act of 2003	2003	−1.95	−2.49

Notes: "MMO" refers to the aggregate AMTR as calculated by Mertens and Montiel Olea (2018). "FF" refers to the aggregate AMTR based on authors' calculations.

Table 2—Estimated Impact of Selected Tax Reforms by Age

	Impact on AMTRs (percentage points)					
	16–24 (1)	25–34 (2)	35–44 (3)	45–54 (4)	55–64 (5)	65+ (6)
Revenue Act of 1964	−1.95	−2.07	−2.30	−2.40	−2.39	−2.34
Revenue Act of 1978	−2.01	−1.90	−1.97	−1.83	−1.87	−1.91
Economic Recovery Tax Act 1981	−0.25	−0.35	−0.39	−0.41	−0.37	−0.26
Tax Reform Act of 1986	−2.37	−4.08	−4.66	−4.56	−4.37	−2.92
Omnibus Budget Reconciliation Act of 1990	0.05	0.33	0.55	0.57	0.46	0.42
Omnibus Budget Reconciliation Act of 1993	0.02	0.10	0.13	0.14	0.13	0.17
Jobs and Growth Tax Relief Reconciliation Act of 2003	−1.12	−2.50	−2.66	−2.62	−2.51	−2.28

issue that arises with this type of calculations is the indexing of the federal tax system starting in 1985.[5] To address this concern, we rescale incomes by the automatic adjustments in bracket widths embedded in the federal tax code.

Table 2 reveals that the tax reforms had a similar impact on AMTRs across all age groups. The sign of the changes in the AMTRs induced by the selected tax reforms is the same for all age groups, such that, say, a tax cut in the aggregate AMTR is indeed a tax cut for individuals of all ages. But the magnitude of these changes varies by age, with the 16–24 and 65+ age groups usually experiencing smaller changes.

based on the Statistics of Income (SOI) from the IRS. This difference in data sources might explain the discrepancies between ours and Mertens and Montiel Olea's numbers.

[5] The Economic Recovery Tax Act (ERTA) of 1981 ruled for automatically increasing personal exemptions, standard deductions, and bracket widths by the percentage change in the CPI starting in 1985.

B. *Specification and Estimation*

First introduced by Sims (1980), SVARs have been widely used to study the joint dynamic behavior of multiple aggregate time series by allowing for general feedback mechanisms. Specifically, SVARs first isolate unpredictable variation in policy and outcome variables and then sort out the contemporaneous causal relationships by imposing identifying restrictions. Since the system allows for all possible dynamic causal effects, any linear (or linearized) dynamic stochastic general equilibrium (DSGE) model can be expressed in a state space form that yields a VAR representation for observables that are available to the econometrician (see Fernández-Villaverde et al. 2007). In addition, SVARs also identify the expected future path of policy variables. This is important for interpreting the estimates, as expectations about the persistence of policy actions are arguably key drivers of the economy's response to discretionary tax changes.

Specification.—The baseline reduced-form VAR specification is

$$\begin{bmatrix} AMTR_t \\ URATE_t \\ PRATE_t \\ \mathbf{X}_t \end{bmatrix} = d + A(L) \begin{bmatrix} AMTR_{t-1} \\ URATE_{t-1} \\ PRATE_{t-1} \\ \mathbf{X}_{t-1} \end{bmatrix} + \begin{bmatrix} e_t^{AMTR} \\ e_t^{URATE} \\ e_t^{PRATE} \\ e_t^{\mathbf{X}} \end{bmatrix}, \tag{1}$$

where d is a constant and $A(L)$ is a $p - 1$ lag polynomial. Here, $p = 2$ is the VAR lag length. We estimate the VAR in equation (1) for each age group separately: $AMTR_t$, $URATE_t$, and $PRATE_t$ are age-specific AMTRs, unemployment rate, and participation rate in year t, respectively, and $\mathbf{X}_t$ is a vector of aggregate control variables. We consider a sample of annual observations for the period 1961–2012.

Variables in $\mathbf{X}_t$ include the log of real GDP per capita, the log of the S&P index, and the federal funds rate, which allows us to capture business cycle dynamics, the monetary policy stance, as well as the effects of bracket creep. To explicitly allow for the feedback from debt to taxes and spending, the log of real government spending per capita (purchases and net transfers), the average tax rate, and the change in log of real federal government debt per capita are also included; given the government's budget constraint, any change in tax rates must eventually lead to adjustments in other fiscal instruments.[6] Further, since tax changes are often motivated by concerns about government deficits and debt accumulation, the inclusion of a set of contemporaneous and past fiscal variables most likely provides relevant information to isolate the unanticipated innovations in tax rates.

However, difficulties may still arise to the extent that some of the tax changes classified as exogenous according to the Romer and Romer's (2010) classification are indeed due to population aging. We think that this type of concern does not apply to our exercise. According to the narrative records of post-war tax policy (see Romer and Romer 2009), legislated tax changes are driven by considerations that are unrelated to population aging. Several other variables could enter the VAR. However, we

[6] Burnside, Eichenbaum, and Fisher (2004) argue that the effects of shocks to government purchases may differ depending on the endogenous response of other fiscal instruments.

note that omitted variables that are orthogonal to the fiscal variables (once lagged business cycle indicators are included in the VAR specification) would not bias the estimated effects of changes in AMTRs.

Estimation.—If the system in equation (1) generates unpredictable innovations to the vector of observables $\mathbf{Y}_t$, then the vector of such reduced-form innovations is a linear transformation of the underlying structural shocks

$$\epsilon_t \equiv \left[\epsilon_t^{AMTR}, \epsilon_t^{URATE}, \epsilon_t^{PRATE}, \epsilon_t^{\mathbf{X}}\right]', \tag{2}$$

such that: (i) $E\left[\epsilon_t\right] = 0$, (ii) $E\left[\epsilon_t \epsilon_t'\right] = \Sigma_\epsilon$ is a diagonal matrix (we further impose $\Sigma_\epsilon \equiv I$, where I is the identity matrix), and (iii) $E\left[\epsilon_t \epsilon_{t-j}'\right] = 0$ for $j \neq 0$. The vector of such structural shocks consists then of exogenous innovations in tax rates and other observables that are uncorrelated with each other. In the SVARs literature, the structural shocks ϵ_t are treated as latent variables that are estimated based on the prediction errors of the observables, $\mathbf{Y}_t$, conditional on the informational content in finite distributed lags of $\mathbf{Y}_t$, that is, $\mathcal{Y}_t \equiv \left[\mathbf{Y}_{t-1}', \ldots, \mathbf{Y}_{t-p}'\right]'$.

We posit that $e_t = \mathcal{H}\epsilon_t$, where $\mathcal{H}$ is a matrix of parameters that determines the impact response of the vector of observables, $\mathbf{Y}_t$, to the structural shocks, ϵ_t, we aim to identify. Specifically, we are interested in identifying the parameters in the first column of $\mathcal{H}$, that is, $\mathcal{H}^{i,1}$, with $i = 1, \ldots, \dim(\mathbf{Y}_t)$, that determine the impact response of the observables, $\mathbf{Y}_t$, to the shock to AMTR, ϵ_t^{AMTR}. Identification of $\mathcal{H}^{i,1}$ is achieved by imposing identifying restrictions and hinges on the availability of a proxy variable, m_t, for the latent structural shock to the tax rate, ϵ_t^{AMTR}, that jointly satisfies the identifying assumptions $E\left[m_t \epsilon_t^{AMTR}\right] \neq 0$ and $E\left[m_t \epsilon_t^{i}\right] = 0$ for $i \geq 2$ (the superscript "i" denotes the ith element of the vector). The first orthogonality condition requires the proxy to be contemporaneously correlated with the underlying shock to the average marginal tax rate. The second orthogonality condition requires the proxy to be contemporaneously uncorrelated with all other structural shocks.

In our case, the age-specific proxies for exogenous tax changes are indicator variables that take on nonzero values at the time of an exogenous tax reform, and zero otherwise. The nonzero values are listed in column 1 through 6 of Table 2. When we estimate the proxy SVARs with aggregate variables only, the nonzero values are those in column 2 of Table 1.

Once the contemporaneous (or impact response) parameters are identified and estimated, the effects of a tax shock in subsequent years is traced out using the estimated system in quation (1). The resulting impulse response functions (IRFs) measure the expected dynamic adjustment of the endogenous variables to the initial shock to the AMTR.

IV. Aggregate Effects of Tax Shocks

In this section we establish new facts on the dynamic response of the US labor market to unanticipated changes in tax rates. Notably, we consider two variables that are key indicators of the state of the labor market (i.e., the aggregate unemployment and participation rate).

To interpret the aggregate results, it is useful to consider the following decomposition of the employment to population ratio:

$$\frac{employment}{population} = \left(1 - \frac{unemployment}{labor\,force}\right) \times \frac{labor\,force}{population}.$$

This decomposition shows that employment as a fraction of the population of 16 years and older is equal to the employment rate (fraction of employed workers in the labor force, one minus the unemployment rate) times the participation rate. Hence, the response of the employment to population ratio to tax cuts is accounted by the response of either the unemployment rate or participation rate, or both.

Here we show that the unemployment rate is indeed quite responsive to tax cuts, whereas the participation rate is not. These results indicate that the response of the employment to population ratio is in fact the mirror image of the response of the aggregate unemployment rate. Unemployment, as opposed to participation, is the key margin for understanding the aggregate response of the labor market to unanticipated and temporary changes in AMTRs.

A. *Aggregate Unemployment Response*

Panel A of Figure 3 shows that the estimated response of the AMTR to the tax shock is highly persistent. Specifically, the AMTR remains below average up to five years after the shock. This high persistence contrasts with the relatively fast mean reversion observed for average personal income tax rates, as estimated by Mertens and Ravn (2013), but in line with Mertens and Ravn (2013), where they estimate the response of real GDP per capita to an AMTR cut.

Panel B of Figure 3 shows the response of the unemployment rate to a 1 percentage point cut in AMTRs. The estimates point to large aggregate effects of tax changes. Notably, the peak response, that occurs one year after the initial shock, implies that a 1 percentage point cut in AMTRs leads to an approximately 0.7 percentage points decrease in the aggregate unemployment rate. In terms of the US labor force (employed plus unemployed) in 2007, a decline of 0.7 percentage points in the unemployment rate amounts to nearly 1.1 million jobs, that are either preserved or created.

We note that a 1 percentage point cut in AMTRs is not an unusual event as the standard deviation of the AMTR is 4.1 percent. These findings, then, point to highly significant and quantitatively large effects of tax changes on aggregate unemployment. The magnitude of these effects is somewhat greater than that found by Mertens and Ravn (2013) in response to exogenous changes in average effective tax rates on personal income. This observation suggests that unanticipated changes in marginal tax rates have larger effects than changes in average tax rates as they operate through incentive effects on intertemporal substitution.

B. *Aggregate Participation Response*

Panel C of Figure 3 shows the response of the aggregate participation rate to an equally-sized 1 percentage point cut in AMTRs. In contrast with the results for

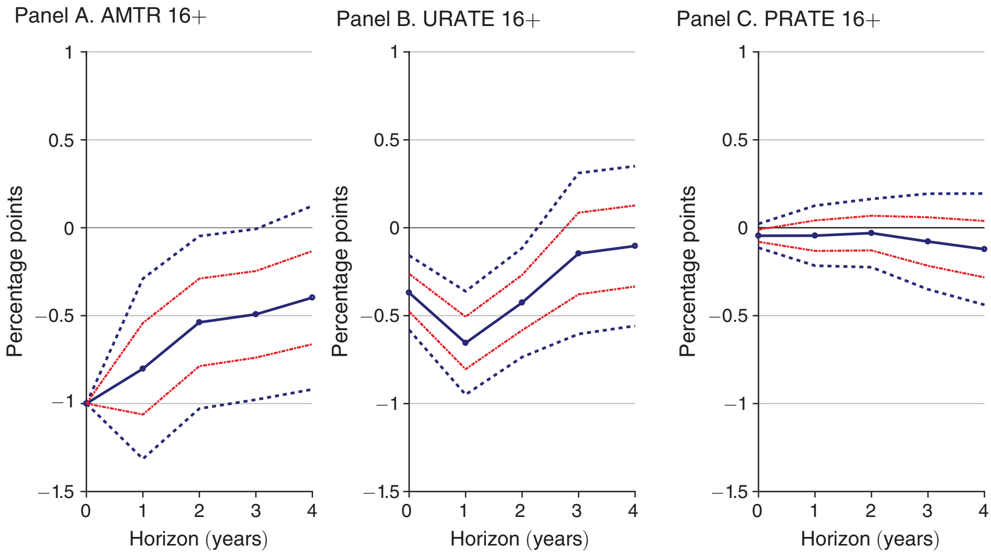

Figure 3. Unemployment and Participation Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal tax rate (AMTR). Full lines with circles are point estimates; dash-dotted lines are 68 percent confidence bands; dashed lines are 95 percent confidence bands. Both intervals are computed using the Delta-method suggested by Montiel Olea, Stock, and Watson (2017) with a Newey and West (1987) HAC-robust residual covariance matrix.

the unemployment rate, the response of the participation rate is both statistically and economically insignificant. This finding is, perhaps, not surprising since the labor force participation rate has displayed pronounced low-frequency movements over the post-war period, arguably driven by long-run demographic trends that are hardly affected by the magnitude of the realized tax shocks. Yet, one cannot a priori dismiss the hypothesis that larger shocks to AMTRs could generate a substantially different response in labor force participation. In that case, the main concern would be whether linear SVARs remain reliable in recovering the true dynamic response to large shocks.

V. Age-Specific Effects of Tax Shocks

In this section we detail the demographics of the US labor market response to tax shocks. To this aim, we study if, and the extent to which, the dynamic responses of unemployment and participation rates to tax cuts vary by age. In doing so, it is imperative to keep in mind that the effects of a shock to the AMTR on age-specific labor-market outcomes incorporate equilibrium feedback effects that result from the fact that changes in average marginal tax rates impact workers in all age groups, rather than just the specific age group considered.

Ultimately, we are interested in decomposing the response of the aggregate unemployment rate into the relative contribution of each age group, after all the equilibrium feedback effects have played out. This will allow us to quantify the relative importance of the young, prime-age, and old in shaping the aggregate response of the unemployment rate to tax cuts.

A. *Age-Specific Unemployment Response*

We now turn to study the role of age-specific unemployment rates and labor force shares for the response of the aggregate unemployment rate. Towards this goal, we consider the following decomposition of the aggregate unemployment rate:

$$\frac{\textit{unemployment}}{\textit{labor force}} = \sum_a \underbrace{\frac{\textit{labor force}_a}{\textit{labor force}}}_{\substack{\text{age-specific}\\ \text{labor force share}}} \times \underbrace{\frac{\textit{unemployment}_a}{\textit{labor force}_a}}_{\substack{\text{age-specific}\\ \text{unemployment rate}}}, \tag{3}$$

where a indicates age and the labor force is defined as employed plus unemployed workers of 16 years and older, in accord with the definition used by the Bureau of Labor Statistics (BLS). The decomposition in equation (3) shows that the response of the aggregate unemployment rate to tax cuts is accounted by the response of either age-specific labor force shares or age-specific unemployment rates, or both. We show that age-specific unemployment rates are indeed responsive to tax cuts, whereas age-specific labor force shares are not.

Labor Force Shares versus Unemployment Rates by Age.—To disentangle the relative contribution of age-specific labor force shares from that of age-specific unemployment rates, we construct a counterfactual time series of the aggregate unemployment rate, u_t^{FLFS}, in which age-specific labor force shares are fixed at their sample averages, $\bar{\phi}_a^{LF}$, whereas age-specific unemployment rates, $u_{a,t}$, vary over time as in the data:

$$u_t^{FLFS} \equiv \sum_a \bar{\phi}_a^{LF} \times u_{a,t}, \quad \text{with} \quad \sum_a \bar{\phi}_a^{LF} = 1. \tag{4}$$

We then reestimate the proxy SVARs by replacing the actual unemployment rate with the counterfactual unemployment rate in equation (4).

Panel B of Figure 4 shows that the impulse response to the AMTR shock of the counterfactual unemployment rate (dashed line with diamonds) is nearly indistinguishable from the impulse response of the actual unemployment rate (full line with circles). We conclude that age-specific unemployment rates, as opposed to age-specific labor force shares, are responsible for the response of the unemployment rate to a tax shock. As argued before, labor force shares by age display marked low-frequency movements in the post-war period. However, such low-frequency movements are due to the underlying demographic trends that pervade the entire US economy, which are unlikely to be affected by temporary changes in marginal tax rates. The composition of the workforce is largely predetermined by fertility decisions made prior to the observed changes in tax rates.

These findings are important for the scope of this paper as they provide an empirically validated restriction, akin to an orthogonality condition, that will

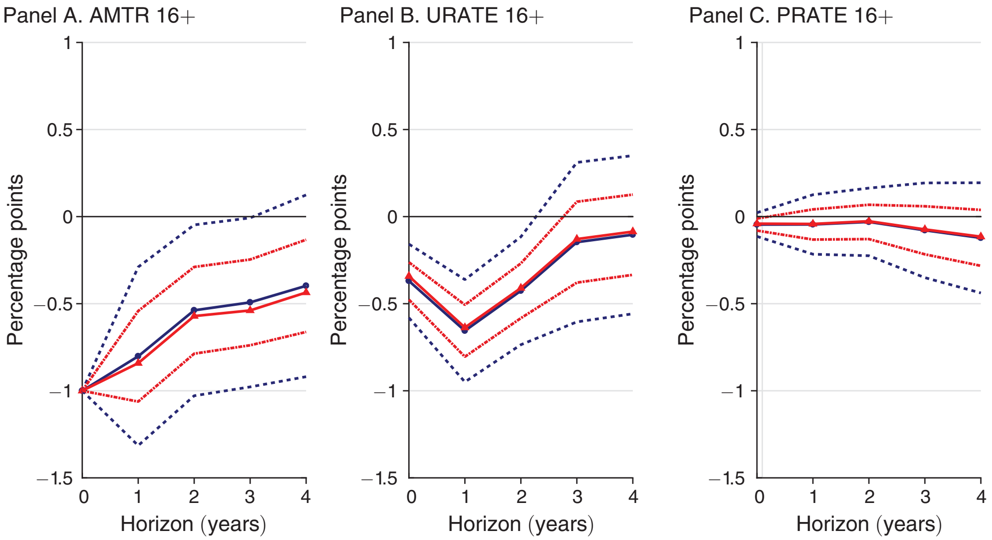

FIGURE 4. ACTUAL DATA VERSUS COUNTERFACTUAL SERIES

Notes: Full lines with circles are point estimates for the response of the actual unemployment rate and participation rate to a 1 percentage point cut in the average marginal tax rate (AMTR); dash-dotted lines are 68 percent confidence bands; dashed lines are 95 percent confidence bands. Both intervals are computed using the Delta-method suggested by Montiel Olea, Stock, and Watson (2017) with a Newey and West (1987) HAC-robust residual covariance matrix. Full lines with diamonds show the response estimated with the counterfactual unemployment and participation rate, as implied by equation (4) and equation (7), respectively.

enable us to quantify the role of an aging labor force in shaping the aggregate unemployment response to tax cuts. Specifically, we can view the IRF of the aggregate unemployment rate as a *weighted* average of the IRFs of the age-specific unemployment rates, where the weights are (sample averages of) the age-specific labor force shares:

$$IRF(h) \approx \sum_a \bar{\phi}_a^{\mathrm{LF}} \times IRF(a,h), \tag{5}$$

where h is the number of years after the initial shock to the AMTR. Note that the impulse response of the counterfactual unemployment rate, shown in panel B of Figure 4, guarantees that the right-hand side of equation (5) is indeed a strikingly good approximation of the impulse response of the actual unemployment rate, the left-hand side of equation (5). With this approximation result at hand, we decompose the contribution of each age group to the aggregate response to tax cuts. We note that if there was no heterogeneity in the IRFs across age groups, then the labor force shares would become irrelevant as $\sum_a \bar{\phi}_a^{LF} = 1$ by construction. Thus, changes in the age composition of the labor force affect the response of the aggregate unemployment rate insofar as the unemployment rate responses to tax cuts differ by age.

Unemployment Rates by Age.—We next establish that the response to tax cuts of the aggregate unemployment rate indeed masks substantial heterogeneity by age. Specifically, the unemployment rate response of the young is nearly twice as large as

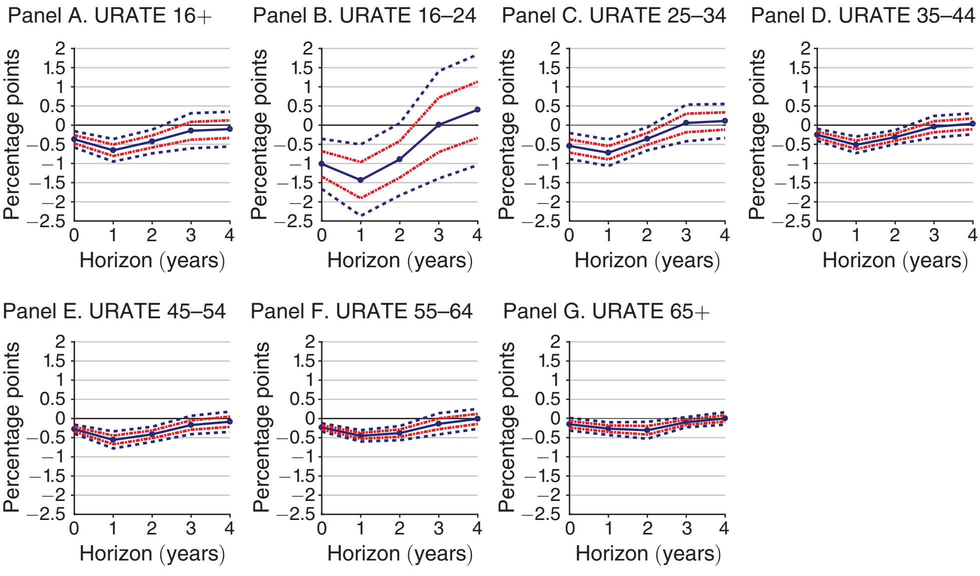

Figure 5. Unemployment Rate Response to a Tax Cut by Age

Notes: The figure shows the response to a 1 percentage point cut in the age-specific average marginal tax rate (AMTR). Full lines with circles are point estimates; dash-dotted lines are 68 percent confidence bands; dashed lines are 95 percent confidence bands. Both intervals are computed using the Delta-method suggested by Montiel Olea, Stock, and Watson (2017) with a Newey and West (1987) HAC-robust residual covariance matrix.

that of the prime-age and old workers in the labor force. This age-specific heterogeneity in the responsiveness to tax cuts is the channel through which shifts in the age composition of the labor force affect the response of the aggregate unemployment rate to tax cuts.

Figure 5 shows the responses of age-specific unemployment rates to an equally-sized 1 percentage point cut in AMTRs. Impulse responses in each panel are obtained using the age-relevant AMTR to obtain the reduced-form residuals and the appropriate instrument to identify age-specific AMTR shocks. The estimates display stark differences in the responses of the young (16–34 age group), prime-age (35–54 age group), and old (55–64 age group). The unemployment rate of the 16–24-year-olds falls by 1.5 percentage points at the peak of the response that occurs one year after the initial shock. The magnitude of the peak response is more than twice as large as that of the aggregate unemployment rate. We report in panel A for the sake of comparison. The unemployment rate of the 25–34-year-olds falls by approximately 0.8 percentage points, which is broadly consistent with the response of the aggregate unemployment rate. The peak responses of the age groups 35–44, 45–54, and 55–64, are instead nearly half as large as those of the 16–24- and 25–34-year-olds.[7] In the online Appendix, Figure OA.5 and OA.6 show that these

[7] Note that these differences in responses remain if, instead of using age-specific AMTRs, one uses aggregate AMTRs. See Figure OA.4 in the online Appendix.

age differences in the unemployment rate responses to the tax shocks are statistically significant.

B. *Age-Specific Participation Response*

We now turn to analyze the role of age-specific participation rates and population shares for the response of the aggregate participation rate. In Section 4, we argued that participation does not respond to tax shocks. Again, this lack of responsiveness may mask heterogeneity by age. To address this concern, we consider the following decomposition of the aggregate participation rate:

$$\frac{\textit{labor force}}{\textit{population}} = \sum_a \underbrace{\frac{\textit{population}_a}{\textit{population}}}_{\substack{\text{age-specific} \\ \text{population share}}} \times \underbrace{\frac{\textit{labor force}_a}{\textit{population}_a}}_{\substack{\text{age-specific} \\ \text{participation rate}}}. \tag{6}$$

Next, we show that age-specific population shares are in fact unimportant in accounting for the response of the aggregate participation rate. Moreover, the participation rates for all age groups are unresponsive to the identified shocks to the AMTR.

Population Shares versus Participation Rates by Age.—To disentangle the contribution of age-specific population shares from age-specific participation rates, we use a counterfactual series of the aggregate participation rate, n_t^{FPS}, in which age-specific population shares are fixed at their sample averages, $\bar{\phi}_a^P$, whereas age-specific participation rates, $n_{a,t}$, vary over time as in the data:

$$n_t^{FPS} \equiv \sum_a \bar{\phi}_a^P \times n_{a,t}, \quad \text{with} \quad \sum_a \bar{\phi}_a^P = 1. \tag{7}$$

We estimate the SVARs by replacing the actual unemployment and participation rate with the counterfactual series with fixed labor force and population shares in equations (4) and (7), respectively.

In Figure 4, panel C shows that the impulse response of the counterfactual participation rate (dashed line with diamonds) is indistinguishable from the impulse response of the actual participation rate (full line with circles). Thus, we conclude that the age-specific population shares are indeed irrelevant for the response of the aggregate participation rate to tax shocks.

Participation Rates by Age.—Figure 6 shows IRFs of the age-specific participation rates to an equally-sized 1 percentage point cut in age-specific AMTRs. The responses of all groups, young (16–24 and 25–34), prime-age (35–44 and 45–54) and old (55–64) are not statistically significant.[8]

[8] Similar results hold if one uses the aggregate AMTR. See Figure OA.7 in the online Appendix.

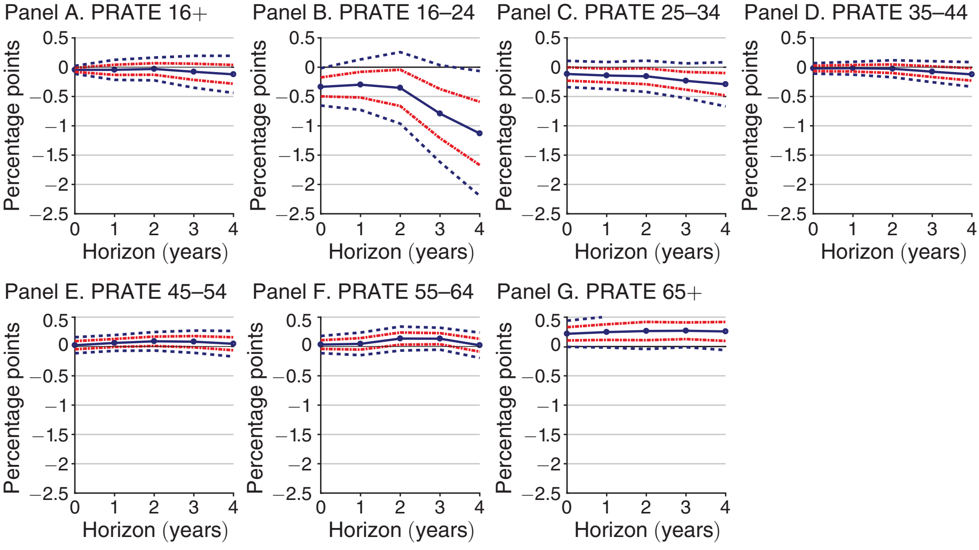

Figure 6. Participation Rate Response to a Tax Cut by Age

Notes: The figure shows the response to a 1 percentage point cut in the age-specific average marginal tax rate (AMTR). Full lines with circles are point estimates; dash-dotted lines are 68 percent confidence bands; dashed lines are 95 percent confidence bands. Both intervals are computed using the Delta-method suggested by Montiel-Olea, Stock, and Watson (2017) with a Newey and West (1987) HAC-robust residual covariance matrix.

VI. Demographic Change and Tax Policy

In this section we build on the empirical evidence in Sections IV and V to quantify the role of demographic change for the effects of tax cuts on the aggregate unemployment rate. We adopt a *quantitative accounting* approach that consists of two steps. First, we measure the relative contribution of each age group to the average response of the unemployment rate to tax cuts. Second, we quantify by how much the aggregate impact of a tax cut changes when we account for the shifts in the age composition of the labor force observed in the data.

A. *Quantifying the Role of Age Composition*

We now turn to quantify the contribution of each age group to the response of the aggregate unemployment rate to tax shocks. To this aim, we decompose the impulse response of the aggregate unemployment rate into additive shares that measure the relative contribution of each group to the aggregate response.

Contribution to the Aggregate Response by Age.—Using the approximation in equation (5), each age group accounts for $share_h^a$ of the response of the aggregate unemployment rate h years after the shock:

$$share_h^a \equiv \bar{\phi}_a^{LF} \times \frac{IRF(a,h)}{IRF(h)}, \quad \text{with} \quad \sum_a share_h^a = 1. \tag{8}$$

Table 3—Shares of Unemployment Response by Age

h (years after shock)	0	1	2	3
$share_h^{16\text{–}24}$	40.06	36.09	35.52	−3.35
$share_h^{25\text{–}34}$	28.18	23.66	18.59	−28.18
$share_h^{35\text{–}44}$	12.62	16.26	15.73	20.08
$share_h^{45\text{–}54}$	11.88	15.11	17.66	69.42
$share_h^{55\text{–}64}$	6.08	7.53	10.06	34.80
$share_h^{65+}$	1.18	1.32	2.44	7.23

Notes: See equation (8) for the definition of $share_h^a$. Shares are reported in percent such that $\sum_{a=16}^{65+} share_h^a = 100$.

Table 3 shows that young individuals, identified by the age group 16–34, account for about two-thirds of the aggregate response on impact and one year after the shock. At the same horizon, the age groups 16–34 and 35–54 account for nearly 93 percent of the aggregate response on impact, and for 91 percent of the response a year after the shock. The combined age groups 55–64 and 65+ account for the remaining 7 percent of the impact response and 9 percent of the lagged response.

How much of the age differences in $share_h^a$ is due to differences in labor force shares versus differences in the ratio of IRF? To answer this question, we use an alternative decomposition in which the labor force shares are set to be the same across age groups, so that $\bar{\phi}_a^{LF} = 1/n_a$, where n_a is the number of age groups:

$$share_h^{a,FLFS} \equiv \frac{1}{n_a} \times \frac{IRF(a,h)}{IRF(h)}. \tag{9}$$

Note that now any age-specific heterogeneity in $share_h^{a,FLFS}$ comes exclusively from the heterogeneity in the unemployment rate responses to tax cuts across different age groups. Thus, the difference between the shares with and without fixed labor force shares can be entirely attributed to the age composition of the labor force.

Indeed, Table 4 points to a quantitatively important role of age composition. Specifically, the two extended age groups 16–34 and 35–54 now account for approximately 83 percent of the aggregate unemployment response on impact and 82 percent of the response a year after the shock. These figures are considerably smaller than those in Table 3. Age composition alone is responsible for a decrease of nearly 10 percentage points. Most of these differences are due to the changes in the relative shares of the age groups 25–34 and 35–44 that represent, on average, roughly 47 percent of the US labor force for the period 1961–2012.

Unemployment Elasticities by Age.—We now study the extent to which the age-specific heterogeneity in the unemployment rate responsiveness to a tax shock can be attributed to differences in unemployment elasticities across age groups versus differences in average unemployment rates. Towards this goal, we use a slightly modified version of equation (8):

$$share_h^a \equiv \bar{\phi}_a^{LF} \times \frac{\bar{u}_a}{\bar{u}} \times \frac{\epsilon_{a,h}^u}{\epsilon_h^u}, \tag{10}$$

Table 4—Counterfactual Shares of Unemployment Response by Age

h (years after shock)	0	1	2	3
$share_h^{16\text{–}24,FLFS}$	40.96	36.39	33.48	−2.31
$share_h^{25\text{–}34,FLFS}$	22.05	18.25	13.41	−14.92
$share_h^{35\text{–}44,FLFS}$	10.25	13.02	11.78	11.03
$share_h^{45\text{–}54,FLFS}$	11.29	14.15	15.46	44.65
$share_h^{55\text{–}64,FLFS}$	9.39	11.46	14.31	36.37
$share_h^{65+,FLFS}$	6.06	6.71	11.53	25.18

Notes: See equation (9) for the definition of $share_h^{a,FLFS}$. Shares are reported in percent such that $\sum_{a=16}^{65+} share_h^{a,FLFS} = 100$.

where $\bar{u}_a$ and $\bar{u}$ are averages of age-specific and aggregate unemployment rates, respectively. To estimate aggregate, ϵ_h^u, and age-specific unemployment rate elasticities, $\epsilon_{a,h}^u$, we estimate the proxy SVARs using the AMTR and unemployment and participation rates in logs, with aggregate variables and separately for each age group. According to equation (10), the share of each group equals the ratio of the age-specific to the aggregate unemployment rate elasticity, weighted by the product of (i) the age-specific labor force share and (ii) the ratio of the age-specific to the aggregate average unemployment rate.

A well-known pattern is that average unemployment rates decrease monotonically with age. This fact has been successfully explained by theories of worker turnover and lifecycle unemployment (see Jovanovic 1979; Chéron, Hairault, and Langot 2013; Esteban-Pretel and Fujimoto 2014; Papageorgiou 2014; Gervais et al. 2016; Menzio, Telyukova, and Visschers 2016). As shown in Table 5, these age differences are large.

Table 6 establishes new facts on the lifecycle profile of the unemployment rate elasticities to tax shocks. The estimates provide evidence of substantial age heterogeneity. Notably, the age group 16–34 features an impact elasticity that is nearly twice as large as that of the age group 55 years and older and 30 percent higher than the age group 35–54.

In our view, these age differences in unemployment elasticities can be instrumental in disciplining theories of lifecycle unemployment, as they provide overidentifying restrictions for quantitative analysis of taxes and unemployment. We further stress that these empirical results also complement the well-known observation that business cycle volatility in labor-market outcomes declines with age (see Clark and Summers 1981; Gomme et al. 2005; Jaimovich and Siu 2009; Jaimovich, Pruitt, and Siu 2013).

B. *Quantifying the Role of Demographic Change*

We now turn to evaluate the quantitative implications of the aging of the baby boomers for the propagation of tax cuts in the United States. Our approach builds on Shimer (1999) and Jaimovich and Siu (2009), where the authors quantify the role

Table 5—Unemployment Rates and Labor Force Shares by Age

Age group	16+	16–24	25–34	35–44	45–54	55–64	65+
Average UNR	6.09	12.48	5.82	4.41	3.91	3.78	3.71
Average UNR-ratio	1	2.04	0.95	0.72	0.64	0.62	0.61
Average LFS	100	18.17	23.74	22.87	19.56	12.04	3.62

Notes: Average unemployment rate (UNR) and labor force share (LFS) for 1961–2012 are reported in percent. The second row indicates average unemployment rates by age, relative to that of 16+-year-olds (UNR-ratio).

Table 6—Unemployment Elasticities by Age

Age group	16+	16–24	25–34	35–44	45–54	55–64	65+
Impact elasticity	1.52	2.14	2.63	1.49	1.79	1.47	1.24
Lagged elasticity	2.79	2.59	3.42	3.00	3.68	3.13	1.92

Notes: "Impact elasticity" refers to the unemployment rate elasticity of a specific age group at horizon $h = 0$; "Lagged elasticity" refers to the unemployment rate elasticity of a specific age group at horizon $h = 1$ (one year after the shock). Each of the impact and lagged elasticities estimates are based on a separate SVARs system that includes the log of the unemployment and the participation rate of a specific age group and a common set of regressors as specified in equation (1). Elasticities are with respect to each age-specific average marginal tax rate (AMTR) and reported in percent. The impact elasticities for the 16–34 and 35–54 age groups are 2.16 and 1.66, respectively. The lagged elasticities for the same age groups are 2.97 and 2.06, respectively.

of age composition of the US workforce for the low-frequency movements in the aggregate unemployment rate and cyclical volatility in hours worked, respectively.[9]

Age Composition-Adjusted Unemployment Response.—Here we show that varying age composition of the labor force has quantitatively important implications for the effectiveness of tax changes. To establish this result, we implement a quantitative accounting exercise. Specifically, we reconstruct the US history of aggregate unemployment responses to a tax cut, by using age-specific labor force shares and unemployment rates observed at a specific point in time, and the unemployment rate elasticities estimated over the sample period 1961–2012, as shown in Table 6. The implied unemployment responses provide a quantitative accounting of how the trends in the age composition of the labor force affect the response of the aggregate unemployment rate to tax cuts.

We construct an aggregate unemployment response to tax cuts, that is adjusted for age composition (AC-adj), as follows:

$$du_{h,t}^{AC\text{-}adj} \equiv \sum_{a=16}^{65+} \phi_{a,t}^{LF} \times \bar{u}_{a,t} \times \epsilon_{a,h}^{u}, \tag{11}$$

[9] The results in Shimer (1999) indicate that the observed trends in the age composition of the labor force have an impact on the *level* of the aggregate US unemployment rate. The entry of the baby boomers in the labor force in the late 1970s, and their aging, accounts for a substantial fraction of the rise and fall in unemployment rates observed in the past 50 years. Jaimovich and Siu (2009) argue that the age composition of the labor force has a causal impact on the *volatility* of hours worked over the business cycle. Since young workers feature less volatile hours worked than prime-age, the aging of the labor force accounts for a significant fraction of the decrease in business cycle volatility observed since the mid-1980s in the United States over the so-called Great Moderation.

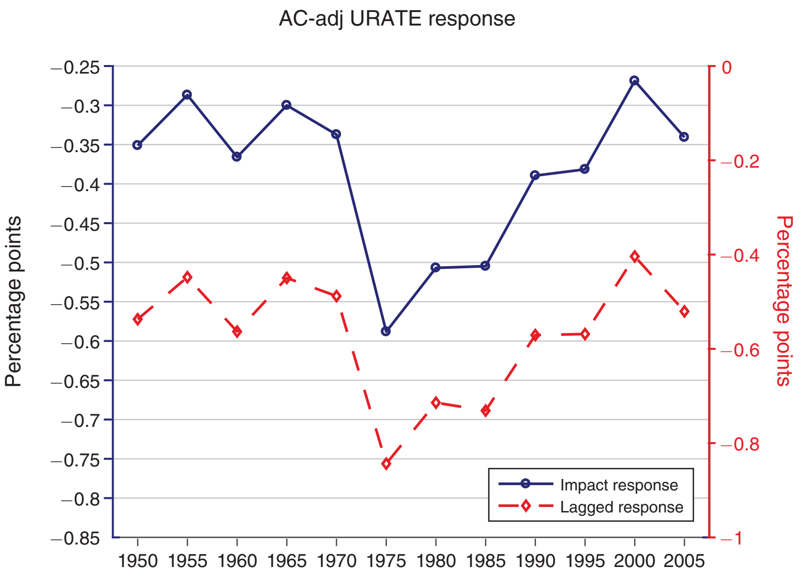

Figure 7. Demographic Change and Unemployment Rate Response to a Tax Cut

Notes: The figure shows the AC-adjusted responses of the unemployment rate of the 16+ age group to a 1 percentage point cut in the average marginal tax rate (AMTR) for all age groups. "Impact response," full line with circles (left axis), shows the AC-adjusted response of the unemployment rate at horizon $h = 0$; "Lagged response," dashed line with diamonds (right axis), shows the AC-adjusted response of the unemployment rate at horizon $h = 1$ (one year after the shock). AC-adjusted responses are constructed using equation (11).

where the time subscripts indicate that the AC-adj response may vary over time due to the observed changes in the age composition of the labor force. Specifically, we use equation (11) to generate unemployment responses to an across-the-board 1 percentage point tax cut, at five-year intervals, from 1950 to 2005.

Before discussing the results, it is worth noticing that this exercise assumes the absence of indirect effects from tax cuts to the age composition of the labor force. This assumption is empirically verified by the results in Sections IV and V (i.e., labor force shares are unresponsive to tax cuts). The responses in equation (11) then measure what would have been the response of the aggregate unemployment rate to a tax cut, if the age composition was that of a specific year in the sample. Hence, given the change in the age composition of the US labor force over the post-war period, the aggregate unemployment response to tax cuts varies over time.

Figure 7 shows the AC-adjusted responses of the aggregate US unemployment rate on impact, and the response a year after the shock. Same patterns hold up to four years after the shock. Note that larger negative values indicate that an equally-sized 1 percentage point cut in the AMTR reduces aggregate unemployment by more. The results indicate large changes in the response of the aggregate unemployment rate to tax cuts. The largest impact and lagged responses occur in the mid-1970s. The peak impact response to a 1 percentage point cut is 0.59 percentage points, which is 42 percent larger than the response of 0.34 percentage points in 2005, whereas the peak lagged response is 0.84 percentage points, which is

38 percent higher than the 0.52 percentage points in 2005. Overall, the magnitude of the counterfactual responses is nearly constant from the mid-1950s to the early 1970s, it markedly increases (in absolute value) during the mid-1970s, and starts to steadily decrease thereafter.

Importantly, this time variation in aggregate unemployment responsiveness lines up nicely with the baby boom and baby bust phenomena we described in Section I. The entry of the baby boomers in the labor force in the 1970s led to a nearly 10 percentage points increase in the share of the 20–34-year-olds. Since the young are more responsive to tax cuts than prime-age and old workers, the responsiveness of the aggregate unemployment rate dramatically increased over that period. However, as the aging of the baby boomers unfolds, the effects of tax cuts on the aggregate unemployment rate are reduced to a level comparable to that of the early 1950s.

Age Composition and Unemployment Rate Elasticity.—The maintained assumption in the counterfactual unemployment rate response in equation (11) is that the unemployment rate elasticities vary by age groups, but they are constant over time. Of course, one can envision the possibility that age-specific elasticities may have changed over time to exactly offset the varying age composition, so that the elasticity of the aggregate unemployment rate remained constant. One way to address this concern is to demonstrate that the estimated aggregate unemployment rate response varies with changes in the age composition. To this aim, the natural approach is to estimate the aggregate regressions in two different subperiods that differ in terms of the age composition of the labor force.

In implementing this approach, we face two challenges. First, tax shocks are not evenly spaced across time such that cutting the sample in subperiods can compromise identification. Second, the labor force shares of the young and prime-age have hump-shaped and U-shaped patterns in the data. This further complicates sample selection as we cannot readily cut the sample in two and estimate regressions separately in each subsample.

Given these difficulties, our strategy is as follows. First, we estimate SVARs including the aggregate unemployment rate and AMTR for 1970–2002, a period characterized by a larger-than-average share of the young in the labor force. Second, we compare the estimates for the subperiod 1970–2002 with those obtained in the full sample 1961–2012. The estimated impact response of the unemployment rate for the period 1970–2002 is 0.69 percentage points, that is, 60 percent higher than the 0.37 estimate in the full sample. The lagged response is 1.18 percentage points compared with 0.65 for the full sample. In terms of elasticities, we estimate an unemployment rate elasticity of 2.13 on impact in the subperiod 1970–2002 that is 34 percent larger than the 1.52 estimate obtained in the full sample. The estimate of the lagged elasticity for 1970–2002 is 5.21, which is roughly 60 percent larger than the 2.79 figure in the full sample. (The same conclusions apply when we consider different subperiods in which we move the beginning and end dates in a five-year window.) Overall, these results confirm our previous accounting that periods with a younger labor force are associated with higher responsiveness of the aggregate unemployment rate to marginal tax rate shocks.

VII. Conclusion

In this paper we investigate the consequences of population aging for the transmission of tax changes to the aggregate labor market in the United States. After isolating exogenous variation in average marginal tax rates in SVARs, we use a narrative identification approach to document that the response of the unemployment rates to tax changes varies significantly across age groups: the unemployment rate response of the young is nearly twice as large as that of the old. This heterogeneity is the channel through which shifts in the age composition of the labor force impact the response of the US unemployment rate to tax changes. We find that the aging of the baby boomers considerably reduces the effects of tax cuts on aggregate unemployment.

These results indicate that the age composition of the labor force is a quantitatively important propagation mechanism of tax policy. Tax changes of the size observed in the United States in the post-war period are becoming increasingly less effective in reducing unemployment. Our estimates suggest that tax cuts targeted toward the young are likely to have larger effects on aggregate unemployment than untargeted ones. Whether age-dependent, countercyclical tax cuts are desirable in terms of macroeconomic stabilization requires the specification of a structural model of the economy. While we view studying the normative implications of our findings as both promising and important, we leave these issues for future research.

Appendix A: Data

A. *Marginal Tax Rates*

This section details how we constructed average marginal tax rates (AMTRs) for 1961–2012. Consistent with Barro and Redlick (2011) and Mertens and Montiel Olea (2018), AMTR (both aggregate and age-specific) is calculated as the sum of the average marginal individual income tax rate (AMIITR) and the average marginal payroll tax rate (AMPTR). Figure A1 and A2 show time series of the aggregate and age-specific AMIITR and AMPTR.

Average Marginal Individual Income Tax Rates.—To calculate AMIITR we follow the procedure in Barro and Redlick (2011) and most of the literature thereafter. AMIITR, both aggregate and by age group, is based on a broad concept of labor income that includes wages, self-employment, partnership, and S-corporation. Our income source is the March supplement of the CPS. The March supplement contains income information (INCWAGE, INCBUS, and INCFARM) and demographic characteristics such as age, marital status, and the number of children in the household. Data are extracted from IPUMS. NBER TAXSIM program simulates marginal tax rates given data inputs on income and demographic characteristics. Aggregate and age-specific AMTRs are calculated as a weighted average of individual AMTRs using adjusted gross income (AGI) as weights. Our aggregate AMTR displays a correlation of 0.92 in levels (and 0.90 in first-difference) with the AMTR calculated by Barro and Redlick (2011) as extended by Mertens and Montiel Olea (2018).

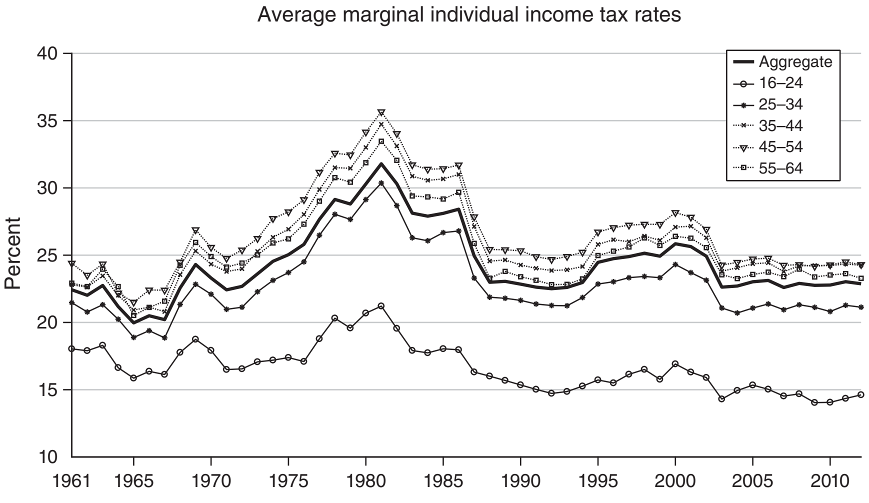

Figure A1. Average Marginal Individual Income Tax Rates, 1961–2012

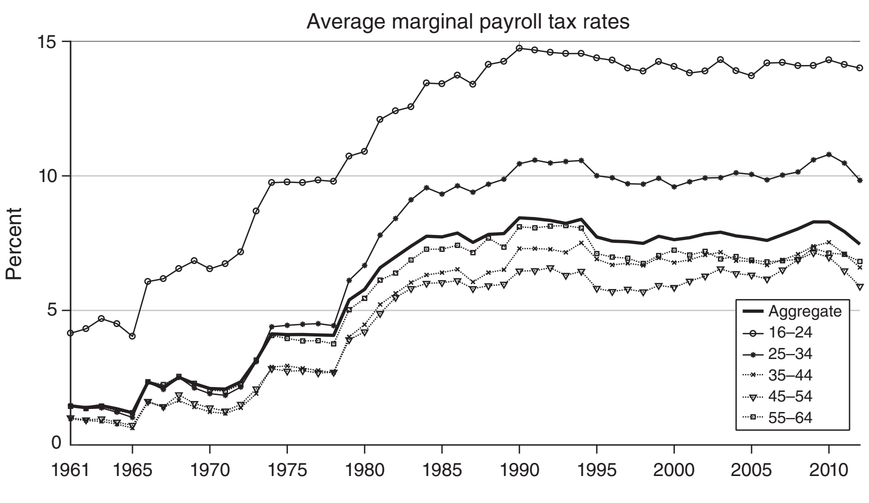

Figure A2. Average Marginal Payroll Tax Rates, 1961–2012

Average Marginal Payroll Tax Rates.—Based on Barro and Sahasakul (1983), AMPTRs are calculated as

$$\text{(A1)} \qquad AMPTR = w_1\left(\frac{s_f + s_w}{1 + s_f}\right) + w_2 s_e,$$

where s_f, s_w, and s_e are the contribution rates paid by employers, employee, and self-employed, respectively, and w_1 and w_2 are total taxable earnings of those with earnings below the annual maximum taxable as a ratio of total income. Data on contribution rates and maximum taxable earnings are available from the Annual Statistical Supplement at http://www.ssa.gov/policy/docs/statcomps/supplement/.

The number of employees and self-employed is obtained from the BLS. The BLS publication "Self-employment In The United States" provides detailed statistics for the aggregate economy and age-specific groups from 1994 onward. To obtain estimates of the years prior to 1994, we use the 1994 observation and impute it backward until the beginning of the sample. Our estimates are robust to alternative imputation choices. For instance, assuming that the share of self-employed in each age group equals the share of self-employed in the aggregate has little impact on the calculations prior 1994. Our aggregate AMPTR displays a 0.97 (0.87) correlation with the level (first-difference) of the AMPTR calculated by Barro and Redlick (2011) as extended by Mertens and Montiel Olea (2018).

B. *Other Time Series*

Real GDP per tax unit is NIPA 1.1.3 line 1 divided by potential tax units. The **Federal Funds Rate** is the annual average effective federal funds rate from the Federal Reserve Board of Governors. **Government Debt** per tax unit is federal debt held by the public, measured by Table L.106 line 19 (federal government, liabilities, credit market instruments) in the US Financial Accounts (release Z.1 of the Federal Reserve Board), divided by the log change in the BLS CPI Research Series Using Current Methods (CPI-U-RS) and potential tax units. **Government Spending** per tax unit is the sum of federal government purchases, net interest rate expenditures and net transfers (NIPA 3.2 line 46 less lines 3, 4, 7, 10, and 11 plus NIPA 3.12U line 25), divided by the CPI-U-RS and potential tax units. The **Real Stock Price** is the S&P composite index from updates of Shiller (2000), divided by the CPI-U-RS. The **Average Tax Rate** is the sum of federal personal current taxes and contributions for social insurance (NIPA 3.2 line 3 plus NIPA 3.7 lines 3 and 21) divided by total market income from Piketty and Saez (2003). The **labor force** is the number of employed plus unemployed persons. The **unemployment rate** is the number unemployed divided by labor force. The **participation rate** is labor force divided by population. **Population** is the civilian noninstitutional population 16 years of age and older. Data for the labor force, unemployment rate, participation rate, numbers of employed and unemployed, and population for the total economy and by age groups are obtained from the CPS, published by the BLS, and available at the CPS home page at http://www.bls.gov/cps/.

REFERENCES

Acconcia, Antonio, Giancarlo Corsetti, and Saverio Simonelli. 2014."Mafia and Public Spending: Evidence on the Fiscal Multiplier from a Quasi-experiment." *American Economic Review* 104 (7): 2185–2209.

Anderson, Emily, Atsushi Inoue, and Barbara Rossi. 2016. "Heterogeneous Consumers and Fiscal Policy Shocks." *Journal of Money, Credit and Banking* 48 (8): 1877–88.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy* 4 (2): 1–27.

Barro, Robert J., and Charles J. Redlick. 2011. "Macroeconomic Effects from Government Purchases and Taxes." *Quarterly Journal of Economics* 126 (1): 51–102.

Barro, Robert J., and Chaipat Sahasakul. 1983. "Measuring the Average Marginal Tax Rate from the Individual Income Tax." *Journal of Business* 56 (4): 419–52.

Brückner, Marcus, and Evi Pappa. 2012. "Fiscal Expansions, Unemployment, and Labor Force Participation: Theory and Evidence." *International Economic Review* 53 (4): 1205–28.

Burnside, Craig, Martin Eichenbaum, and Jonas D.M. Fisher. 2004. "Fiscal Shocks and Their Consequences." *Journal of Economic Theory* 115 (1): 89–117.

Caldara, Dario, and Christophe Kamps. 2017. "The Analytics of SVARs: A Unified Framework to Measure Fiscal Multipliers." *Review of Economic Studies* 84 (3): 1015–40.

Chéron, Arnaud, Jean-Olivier Hairault, and François Langot. 2013. "Life-Cycle Equilibrium Unemployment." *Journal of Labor Economics* 31 (4): 843–82.

Clark, Kim B., and Lawrence H. Summers. 1981. "Demographic Differences in Cyclical Employment Variation." *Journal of Human Resources* 16 (1): 61–79.

Cloyne, James. 2013. "Discretionary Tax Changes and the Macroeconomy: New Narrative Evidence from the United Kingdom." *American Economic Review* 103 (4): 1507–28.

Esteban-Pretel, Julen, and Junichi Fujimoto. 2014. "Life-Cycle Labor Search with Stochastic Match Quality." *International Economic Review* 55 (2): 575–99.

Favero, Carlo, and Francesco Giavazzi. 2012. "Measuring Tax Multipliers: The Narrative Method in Fiscal VARs." *American Economic Journal: Economic Policy* 4 (2): 69–94.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson. 2007. "ABCs (and Ds) of Understanding VARs." *American Economic Review* 97 (3): 1021–26.

Gervais, Martin, Nir Jaimovich, Henry E. Siu, and Yaniv Yedid-Levi. 2016. "What Should I Be When I Grow Up? Occupations and Unemployment over the Life Cycle." *Journal of Monetary Economics* 83: 54–70.

Gomme, Paul, Richard Rogerson, Peter Rupert, and Randall Wright. 2005. "The Business Cycle and the Life Cycle." *NBER Macroeconomics Annual 2004* 19: 415–61.

Hansen, Gary D., and Selahattin İmrohoroğlu. 2009. "Business Cycle Fluctuations and the Life Cycle: How Important Is On-the-job Skill Accumulation?" *Journal of Economic Theory* 144 (6): 2293–2309.

House, Christopher L., and Matthew D. Shapiro. 2006. "Phased-In Tax Cuts and Economic Activity." *American Economic Review* 96 (5): 1835–49.

House, Christopher L., and Matthew D. Shapiro. 2008. "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation." *American Economic Review* 98 (3): 737–68.

Jaimovich, Nir, Seth Pruitt, and Henry E. Siu. 2013. "The Demand for Youth: Explaining Age Differences in the Volatility of Hours." *American Economic Review* 103 (7): 3022–44.

Jaimovich, Nir, and Henry E. Siu. 2009. "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility." *American Economic Review* 99 (3): 804–26.

Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *Journal of Political Economy* 87 (5): 972–90.

Menzio, Guido, Irina A. Telyukova, and Ludo Visschers. 2016. "Directed Search over the Life Cycle." *Review of Economic Dynamics* 19: 38–62.

Mertens, Karel, and José Luis Montiel Olea. 2018. "Marginal Tax Rates and Income: New Time Series Evidence." *Quarterly Journal of Economics* 133 (4): 1803–84.

Mertens, Karel, and Morten O. Ravn. 2012. "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks." *American Economic Journal: Economic Policy* 4 (2): 145–81.

Mertens, Karel, and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review* 103 (4): 1212–47.

Mertens, Karel, and Morten O. Ravn. 2014. "A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers." *Journal of Monetary Economics* 68 (S): S1–19.

Montiel Olea, J.L., J. Stock, and M. Watson. 2017. "Inference in SVARs Identified with External Instruments." Unpublished. https://www.princeton.edu/~mwatson/papers/SVARIV.pdf.

Nakamura, Emi, and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104 (3): 753–92.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–08.

Papageorgiou, Theodore. 2014. "Learning Your Comparative Advantages." *Review of Economic Studies* 81 (3): 1263–95.

Pappa, Evi. 2009. "The Effects of Fiscal Shocks on Employment and the Real Wage." *International Economic Review* 50 (1): 217–44.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics* 118 (1): 1–39.

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

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 71–162. Amsterdam: North-Holland.

Ramey, Valerie A., and Matthew D. Shapiro. 1998. "Costly Capital Reallocation and the Effects of Government Spending." *Carnegie-Rochester Conference Series on Public Policy* 48: 145–94.

Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy* 126 (2): 850–901.

Ríos-Rull, José-Víctor. 1996. "Life-Cycle Economies and Aggregate Fluctuations." *Review of Economic Studies* 63 (3): 465–89.

Romer, Christina D., and David H. Romer. 2009. "A Narrative Analysis of Postwar Tax Changes." http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=AA0401CCFE2C3020F5967419246B79F1?doi=10.1.1.370.4762&rep=rep1&type=pdf.

Romer, Christina D., and David H. Romer. 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review* 100 (3): 763–801.

Shiller, Robert J. 2000. *Irrational Experience.* Princeton, NJ: Princeton University Press.

Shimer, Robert. 1999. "Why Is the U.S. Unemployment Rate So Much Lower?" *NBER Macroeconomics Annual 1998* 13: 11–74.

Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica* 48 (1): 1–48.

Wong, Arlene. 2015. "Population Aging and the Transmission of Monetary Policy to Consumption." https://economics.yale.edu/sites/default/files/files/Events/jf-2016/Arlene_Wong_JMP_Latest.pdf.

This article has been cited by:

1. Henrique S. Basso, Omar Rachedi. 2021. The Young, the Old, and the Government: Demographics and Fiscal Multipliers. *American Economic Journal: Macroeconomics* **13**:4, 110-141. [Abstract] [View PDF article] [PDF with links]