

SFB
823

The employment dynamics of different population groups over the business cycle

Christian Bredemeier, Roland Winkler

Nr. 21/2014

Discussion Paper



The employment dynamics of different population groups over the business cycle

Christian Bredemeier, University of Cologne
Roland Winkler, TU Dortmund University

May 2014

Abstract

We use Bayesian VARs to analyze differences in employment dynamics across population groups in the US. The employment of males, young people, non-whites, the less educated, and workers in blue-collar occupations fluctuates more heavily. We further document across-group differences in the timing of employment reactions to business-cycle shocks. Groups with stronger fluctuations are affected early in the transmission process. Furthermore, we disentangle between the effects of different shocks. Supply shocks primarily drive differences between occupations, genders, age groups, and education groups. Demand shocks are important for the dynamics of employment ratios between races and ethnic origins.

Keywords: employment, business cycles, heterogeneous labor-market outcomes, demography

JEL classification: E24, E32, J10, J21

1 Introduction

Whose jobs are cut in a recession and for whom are jobs created in booms? Who is fired first in a recession and who finds a job early in a boom? In this paper, we seek to answer these questions and aim to identify the drivers of heterogeneous employment dynamics. To this end, we estimate, using Bayesian vector autoregressions, the effects of business-cycle shocks on the employment dynamics of different population groups in the United States.

An examination of the cyclical employment differences across population groups is important because of several reasons. First, an aggregate perspective can hide whether business-cycle costs are asymmetrically borne by certain population groups. Moreover, differences in employment dynamics have impacts on the cyclical component of inequality and, thus, affect also aggregate

Bredemeier: University of Cologne, Center for Macroeconomic Research, Albertus-Magnus-Platz 1, 50923 Cologne, Germany, bredemeier@wiso.uni-koeln.de. Winkler: TU Dortmund University, Vogelpothsweg 87, 44227 Dortmund, Germany, roland.winkler@tu-dortmund.de. Parts of this research have been produced while Bredemeier was at TU Dortmund University. Financial support from German Science Foundation (DFG) through SFB 823 is gratefully acknowledged.

costs of business cycles. Finally, for policymakers that seek to stabilize employment over the cycle, it is important to know where in the economy initial job losses in recessions are located.

In order to examine the employment dynamics of different population groups over the business cycle, we estimate a set of VAR models. We use trend-filtered quarterly US data on labor productivity, real GDP, employment, the real interest rate, and employment of different population subgroups. We focus on the reaction of employment ratios between population subgroups, constructed along the following 6 dimensions: gender, employment status, occupation, age, race/ethnic origin, and educational attainment. As the population-group time-series data is available at rather short sample length, we apply Bayesian estimation techniques that allow us to use information from earlier available aggregate data in an empirical prior.

We identify, using sign restrictions, two business-cycle shocks: a productivity and a non-productivity shock. As sign restrictions, we impose that a positive productivity shock raises labor productivity and real GDP, while a positive non-productivity shock raises real GDP without raising labor productivity. In our estimates, the identified non-productivity shock moves economic activity and real interest rates up and labor productivity down. From a theoretical perspective, such responses are consistent with positive demand-side innovations that are not triggered by monetary policy innovations or, in terms of a textbook model, rightward shifts of the IS curve. Examples are fiscal policy shocks, investment shocks, consumption preference shocks, or shocks to net exports. For this reason, we refer to the non-productivity shock as a "demand shock" and to the productivity shock as a "supply shock".

Using our estimates, we find significantly different employment dynamics of different population groups over the cycle. In booms, predominantly the employment of young people, males, the less educated, ethnic minorities, and employees in construction and production occupations is boosted. This is in line with findings of an existing literature, e.g. Elsby et al. (2010) or Hoynes et al. (2012), which we review below.

What we learn on top of that from our dynamic analysis is that the composition of employment changes particularly at the beginning of the transmission process of business-cycle shocks. Generally, employment ratios are found to be leading aggregate employment. Groups who are particularly strongly affected by business-cycle shocks, for example men and workers in blue-collar

occupations, also tend to be affected early, before the impacts of the shock spread over to women and workers in white-collar occupations.

Disentangling these heterogeneous employment fluctuations to the effects of different shocks, we find that supply and demand shocks have qualitatively similar effects on most groups. However, the quantitative effects differ substantially for many groups. We find that the stronger pro-cyclicality in the employment of males, the young, the less educated, and blue-collar workers is primarily caused by supply shocks. By contrast, cyclical fluctuations in employment ratios between races and ethnic origins are foremost initiated by demand-side innovations. We further find that supply shocks impact relatively early on the composition of employment. By contrast, demand shocks impact later on the employment ratios. The fact that most employment ratios are leading aggregate employment in the data is thus foremost explained by supply-side innovations.

Our results have several interesting implications. A first implication of our results is that business-cycle costs are disproportionately borne by poorer population groups and that inequality rises in recessions. A second implication is a trade-off between the stabilization of aggregate employment and the stabilization of its composition when policy tries to use demand-side impulses to counter negative supply shocks. Third, the results of this paper improve our understanding about the chain of events in a boom or bust. In a supply-side recession, workers in sectors with high shares of male, blue-collar employment are laid off early before the recession spreads over to other parts of the economy.

Potential explanations for our findings can lie in the demographic composition of sectoral employments and the different importance of hiring rates across population groups. Men and the less educated have higher likelihoods to work in construction-intense sectors, the public sector is an important employer for blacks. Dynamics in hiring rates are arguably important for the young. We show that dynamics in sectoral activity and in hiring rates hint to what can lie behind the documented heterogeneous employment dynamics.

Our paper contributes to an existing literature that explores differences in employment cyclicalities across population groups. A first strand of this literature examines which groups' employments are more strongly exposed to the business cycle by either looking at unconditional moments or by regressing groups' employments on the state of the cycle. Examples include Clark and Summers

(1981), Hynes et al. (2001), Jefferson (2008), Elsbey et al. (2010), or Hoynes et al. (2012). Based on a different method, our results in this respect are consistent with their findings.

We make two additions to this literature. First, we point out that employment differences between population groups typically lead aggregate employment. This motivates our study of the whole dynamics of groups' employments over the business cycle. Second, we study the reaction of different groups' employment to the sources of business-cycle fluctuations. This helps us to identify the drivers of the different cyclicalities across population groups.

Considering specific sources of heterogeneous employment variations, a related literature estimates the effects of monetary policy and fiscal policy. Studies that estimate the responses of different groups' employment to monetary policy shocks are, for example, Abell (1991), Zavodny and Zha (2000), Thorbecke (2001), Carpenter and Rodgers (2004), Jefferson (2005), Rodgers (2008), and Williams (2011). We do not investigate monetary policy shocks here. Instead, we consider a broadly defined aggregate demand shock, but one that turns out to exclude monetary policy innovations. One specific demand shock, increases in military spending, has been analyzed with respect to its effects on different population groups by Giavazzi and McMahon (2013). As does our broad demand shock, military spending shocks increase the employment of young individuals in their results. Concerning the gender effects of military spending, Giavazzi and McMahon (2013) report different results than we find for the broad demand shock.

The aforementioned literature focuses on demand-side disturbances. Though, supply shocks are important drivers of business-cycle fluctuations as well. Therefore, we believe it is important to study the employment reactions of different population groups in a empirical model that is driven not only by demand-side but also by supply-side disturbances. Our results suggest that the unconditional cyclicalities in the data are in fact mainly explained by supply-side disturbances.

The paper is comprised of four main sections. Section 2 provides unconditional moments of our employment data. Section 3 describes our estimation strategy. Section 4 presents the estimation results and discusses their implications as well as some potential explanations. Section 5 concludes.

2 Unconditional moments

Before we present our empirical strategy, let us first take a look at the unconditional cyclicalities in our data. The data source for employment is the Labor Force Statistics from the Current Population Survey provided by the Bureau of Labor Statistics (see Table A in the Appendix for details). We compare the cyclical employment behavior of different population groups by analyzing the log ratio between the detrended and seasonally adjusted employments of two groups. We consider a total of 34 employment ratios between 22 population groups built along six dimensions: gender, employment status, occupation, age, race/ethnic origin, and educational attainment.¹ Column 2 of Table 1 lists all 34 employment ratios within the six dimensions.

– insert Table 1 here –

Column 3 of Table 1 reports the contemporaneous correlations of the 34 employment ratios with aggregate employment which we, now and henceforth, take as the indicator of the business cycle. A positive sign indicates that the first group has a more strongly pro-cyclical employment than the second group. As we consider seasonally adjusted, detrended data, this is comparable to the relative cyclicity estimates in, e.g., Hoynes et al. (2012). As a consequence, the unconditional correlations in column 3 of Table 1 mirror their results. The negative entry in row 1 indicates that female employment is less pro-cyclical than male employment. Further, we observe particularly strong pro-cyclicalities of full-time employment, of employment in construction or production occupations, as well as of employment of relatively young workers, of the less educated, and of non-whites.

The contemporaneous correlation gives us only a limited picture of the correlation function between the respective employment ratio $EmpRatio$ and aggregate employment $AggEmp$, that is $Corr(EmpRatio_t, AggEmp_{t-i})$. In Figure 1, we show the correlation functions to aggregate employment for two employment ratios within a ± 8 quarters window. As examples, we choose the female-male employment ratio and the ratio between employment in production occupations and employment in management, professional and related occupations. We see that correlations with leads of aggregate employment ($i < 0$) exceed the contemporaneous correlations in absolute value.

¹Where available, we use employment-to-population ratios. Where this is not available in the data, we work with employment head counts by group. The cyclical behavior of employment ratios between such groups is similar to the ratio in the employment-to-population ratios as long as there is no substantial cyclical component in population size of the different groups considered.

The strongest correlations are found with aggregate employment three quarters ahead ($i = -3$), thus these employment ratios lead aggregate employment.

– insert Figure 1 here –

For all employment ratios considered, we report, in columns 4 and 5 of Table 1, the lag i at which we find the strongest correlation within a ± 8 quarters window, together with the value of this strongest correlation. Notably, we find that almost all employment ratios are leading aggregate employment, with only a few exceptions for which, though, correlations are weak. This implies that groups which are more strongly exposed to employment fluctuations also loose their jobs early in a recession or find jobs early in booms. By contrast, groups with comparably low employment fluctuations seem to have the most pronounced changes in employment closer to business-cycle peaks. Different population groups thus do not only show differently strong employment reactions over the cycle but also seem to face heterogeneous transmission processes. To investigate this further, we analyze the heterogeneous employment responses within a vector autoregressive model which is the standard framework to analyze the transmission of business-cycle shocks empirically.

3 Methodology

For each of 34 pairs of population groups, we estimate a five-variable VAR with a constant term and three lags. The variables we use are: labor productivity defined as log output per hour, log real GDP per capita, the log aggregate employment-to-population ratio, the real interest rate defined as the difference between the federal funds rate and the annualized change in the GDP deflator, and the log employment ratio between the considered population groups (see Tables A to C in the Appendix for details on data sources and data construction). We use employment ratios rather than the different groups' employments separately because a significant reaction of an employment ratio is tantamount to a *significantly* different response of the two population groups to the considered shock. All time series are seasonally adjusted and HP-filtered with $\lambda = 1600$.

The reduced form VAR reads

$$Y_t = A_0 + A_1 Y_{t-1} + A_2 Y_{t-2} + A_3 Y_{t-3} + \Xi_t,$$

where the vector Y_t includes the variables described above and the matrix $A = (A_0, A_1, A_2, A_3)'$

collects the coefficient matrices. Let V and Σ denote the variance-covariance matrices of A and of the vector of reduced-form residuals Ξ_t , respectively.

Our main sample starts in 2000Q1 because since then, the Bureau of Labor Statistics publishes official disaggregated labor-market data at quarterly frequency. This gives us 245 observations which is rather few compared to the 95 parameters to be estimated. We apply Bayesian estimation techniques to overcome this problem. We construct an empirical prior by using information on the dynamics of the four aggregate variables for which longer times series (starting in 1954Q1) are available. The prior Θ for the Bayesian estimation consists of an estimate for the variance-covariance matrix Σ , a prior for the coefficients matrix A (prior means), and a prior for the variance-covariance matrix of A , V (prior uncertainty).

We run three OLS estimations to construct our empirical prior $\Theta = \{\hat{\Sigma}, A_{prior}, V_{prior}\}$. First, we estimate the full VAR with all five variables for the sample 2000Q1-2012Q4. The resulting variance-covariance matrix of the residuals is our estimate $\hat{\Sigma}$.² Second, we estimate a VAR(3) which contains all variables but the employment ratio for the sample 1954Q1-1999Q4. This gives us subset of coefficients for A_{prior} (those that do not refer to the employment ratio). Third, we estimate an AR(3) process using the log employment ratio between the considered population groups for the sample 2000Q1-2012Q4 resulting in another set of entries for A_{prior} (those that relate the employment ratio to its own lags). For the remaining entries, which describe the interaction between the aggregate variables and the employment ratio, our prior is zero (i.e., we do not use prior information on the reaction of the employment ratio to changes in aggregate variables).

Finally, we use the following hyper parameters in the prior variance-covariance matrix V_{prior} . All off-diagonal entries of V_{prior} are zero, i.e. parameters are independent. For constants, we set a prior variance of 100. For coefficients that refer a lag of the same variable we set a prior variance of 0.05 and for coefficients that refer to a lag of another variable, we use a prior variance of 0.025. The relatively tight prior implies that we shrink the reactions of the employment ratios on aggregate variables relatively strongly towards zero. This works against our aim to identify heterogeneous reactions across population groups.

Once we have estimated the reduced-form VAR by combining prior with likelihood, we identify

²As in the original Minnesota prior (e.g. Doan et al., 1984), it is assumed that there is no prior uncertainty regarding Σ .

productivity and non-productivity shocks using sign restrictions. Table 2 summarizes the applied sign restrictions. A productivity shock is identified as changing labor productivity and GDP in the same direction on impact and in the three following quarters (periods 0-3 after the shock). By contrast, a positive non-productivity shock is identified as raising, over the same time horizon, GDP and aggregate employment but not raising labor productivity.

– insert Table 2 here –

Formally, our identification proceeds as outlined in Uhlig (2005). We take a draw from the posterior distribution of the coefficient matrix A and calculate the Cholesky decomposition of the estimated covariance matrix $\hat{\Sigma} = BB'$. We then take a draw ω from the five-dimensional unit sphere by applying a qr-decomposition of a 5×5 matrix of random numbers drawn from the standard normal distribution. We consider shocks $b = B\omega$ and the impulse response functions to b . If they satisfy the sign restrictions, we keep the draw (A, ω) and save the impulse response functions. We repeat this until we have 50,000 responses to both productivity and non-productivity shocks. We then calculate, for each shock, the median responses and the 90, 80, and 68 percent confidence bands, which reflect Bayesian parameter uncertainty and identification uncertainty.

4 Estimation results

This section presents our estimation results. To start with, we discuss results for the four aggregate variables in our VAR. Their impulse response functions to positive productivity and non-productivity shocks are displayed in Figure 2.³ In the figures, the solid line represents the median response and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points. The magnitude of the considered shock is chosen such that the median response of aggregate employment peaks at an absolute value of one percent. The choice of magnitude is taken for the sake of comparison between productivity and non-productivity shocks.

– insert Figure 2 here –

³Every time we are considering an alternative employment ratio, we re-estimate our 5-variables VAR and obtain different impulse responses for all variables. However, it turns out that the estimated responses of output, aggregate employment, the real interest rate, and labor productivity are barely affected by rotating-in a different employment ratio. In Figure 2, we consider results from the VAR with the gender employment ratio. Results from the other VARs are available upon request.

In response to a positive productivity shock (upper panel of Figure 2), we observe a rise in labor productivity and real GDP, as required by the sign restrictions. Furthermore, we see a hump-shaped rise in total employment and an insignificant response of the real interest rate. Note that the sign of the employment reaction is neither imposed by the sign restrictions nor implied by the choice of magnitude of the shock. The lower panel of Figure 2 displays responses to a positive non-productivity shock. As required by the sign restrictions, real GDP and employment rise whereas labor productivity falls. We further observe a significantly positive reaction of the real interest rate. The estimated responses allow to interpret the non-productivity shock as a demand shock. Analogously, we refer to the productivity shock as a supply shock. Since the real interest rate rises in demand-side triggered booms, our demand shock does not include monetary policy innovations (after which the real interest rate would fall). Our demand shock includes, for example, fiscal policy shocks, investment shocks, consumption preference shocks, or shocks to net exports.

We find it noteworthy that the employment reaction to both shocks is very similar. By construction of the considered shock, the quantitative reaction is similar. But, furthermore, both responses are also very similar in shape. Both responses are hump-shaped and peak responses occur after 4 and after 3 quarters, respectively. Neither this nor the sign of the employment response to supply shocks are imposed by the sign restrictions or implied by the magnitude of the considered shock. This similarity facilitates the interpretation of the employment reaction of the various population groups to supply and demand shocks as we are considering shocks that have very similar effects on aggregate employment.

We now turn to the reactions of the employment ratios between the different population groups, which we present in two formats. First, we look at the impulse response functions of the different employment ratios to positive supply and demand shocks. Figures are layouted in the same way as Figure 2. Second, in Table 3, we summarize the statistically significant employment responses to supply and demand shocks. Employment ratios for which we do not find any significant response are excluded from the table. The columns labeled 'signif.' show the significance level of the most significant response, where '*', '**', and '***' stand for 68, 80, or 90 percent significance. The columns 'periods' list the periods after the shock for which the response is at least significant at 68 percent. The columns 'peak' give the value of the impulse response function which is largest

in absolute value and statistically significant at at least 68 percent. The columns 'at $t =$ ' give the period of this strongest significant response.

– insert Table 3 here –

– insert Figure 3 here –

Gender. Figure 3 presents the responses of the female-male employment ratio. We observe a significant and persistent fall in the employment ratio after both shocks. This confirms that predominantly male employment rises in booms and falls in busts as documented by previous studies. Our results indicate that this observation is qualitatively independent of the source of business-cycle fluctuations. However, there is a substantial quantitative difference between the reactions to the two shocks (which have similar effects on aggregate employment). The median peak response of the employment ratio to a demand shock is only about half as strong as the reaction to a supply shock. After a supply shock, men's employment rises by 1.09 percentage point more than women's, while this number is only 0.51 percent after a demand shock, see Table 3.

The reactions of the employment ratio can be translated into reactions of the groups' employment.⁴ Since the male and female population are of roughly the same size, male and female employment are about 1.4% and 0.6% above trend at the peak of aggregate employment after a supply shock (where, at $t = 4$, the trend deviations of aggregate employment and the employment ratio are 1% and -0.8%). Notice that, while the absolute numbers depend on the scaling of the considered shock, the fact that men's reaction is about 2.5 times higher than women's is independent of scaling. By contrast to the supply shock, the demand shock causes less pronounced differences between genders. Here, the male employment reaction exceeds the female by only about factor 5/3 ($+1.25\%$ and $+0.75\%$ parallel to a demand-caused one-percent rise in aggregate employment.)

Concerning the timing of the responses, we find that the effects of the demand shock are more sluggish than those of the supply shock. While the supply shock has its strongest effect after

⁴For this and the following employment ratios, we proceed as follows. Consider two groups which sum up to the total population with population shares w_1 and w_2 . Let, in any given period, \hat{n}_1 , \hat{n}_2 , $\hat{n} = w_1\hat{n}_1 + w_2\hat{n}_2$, and $\hat{r} = \hat{n}_1 - \hat{n}_2$ denote the percentage trend deviations of employment in group 1, employment in group 2, aggregate employment, and of the employment ratio. From our estimations, we take the medians of \hat{n} and \hat{r} and calculate $\hat{n}_2 = \hat{n} - w_1\hat{r}$ and $\hat{n}_1 = \hat{n}_2 + \hat{r}$. For dimensions with more than two groups, we sum up groups to two larger groups (e.g., production and construction vs. other occupations). We then calculate the corresponding employment ratio \hat{r} between the larger groups as the population weighted averages of the ratios between the smaller groups and then proceed as above. To obtain the change in the total number of jobs by group we multiply percentage trend deviations by the long-run means of employment in this group.

one quarter, the peak response to the demand shock occurs after three quarters. So, the supply shock induces a peak response in the employment ratio before the peak response in aggregate employment and thus explains in parts why the employment ratio leads aggregate employment in the unconditional moments, see Table 1. By contrast, after a demand shock, aggregate employment and the employment ratio peak in a synchronized way.

Let us inspect the transmission process after a favorable supply shock in more detail. There is first a phase where predominantly men are hired while the employment of women is still relatively stable. Thus, the employment ratio between these groups reacts rather strongly in this phase but aggregate employment does not do so yet. Our median responses imply that, one quarter after a positive supply shock, when the gender employment ratio peaks, about 50 percent of all new jobs have been already created while this number is already as high as 75 percent for men. After the peak of the employment ratio but before aggregate employment peaks, more women than men are hired. This brings the employment ratio already somewhat back to normal while aggregate employment still rises. After the peak in aggregate employment, the slump begins with a phase where male and aggregate employment already decline but female employment continues to rise. Symmetrically, in a supply-side caused recession, men - who are hit hardest by layoffs - start to find new jobs more rapidly than women in the recovery. This result is consistent with the observation made in the great recession (Hoynes et al., 2012).

– insert Figure 4 here –

Part time / full time. Figure 4 shows that both shocks cause full-time employment relative to part-time employment to rise. The median peak response is only slightly stronger in the case of supply shocks, for which the employment ratio drops by 2.3 percent, compared to 2.18 percent in the case of demand shocks. Notice that both shocks induce a peak response after 4 quarters, thus at the same time or even later than the peak response of aggregate employment, see Figure 2. This indicates that our empirical model with the two identified shock processes does not explain why this particular employment ratio leads aggregate employment. The fact that the quantitative peak response of the employment ratio is higher than the one of aggregate employment points to the importance of flows between part-time and full-time employment over the cycle.

– insert Figure 5 here –

Occupation. Figure 5 shows the responses of the employment ratios built along the occupation dimension. Both supply and demand shocks affect predominantly employees in construction and production occupations. In booms, employment in production and construction rises relative to that in sales and office, management and professionals, or service occupations. Among the two most affected occupations, employment in construction rises relative to employment in production. Regarding the source of the heterogeneous reactions, our results reveal that the main reason seem to be supply shocks which cause both stronger and more significant employment reactions. At the peak of a productivity-induced rise in aggregate employment, our median responses imply that about 68 percent of the new jobs are created in construction and production occupations. This is substantially more than the long-run share of construction and production in total full-time employment which is about 25 percent. In a demand-caused boom, the implied share of new jobs created in construction and production amounts to 43 percent.

It is further interesting to consider the timing of the effects of supply shocks on the occupational composition of employment. Look, e.g., at the relative reactions of construction employment to supply shocks (third line in panel (a) of Figure 5), which are the most statistically and quantitatively significant cases. The responses peak on impact, so four quarters before aggregate employment does. This indicates that the transmission of shocks operates early through the occupational composition of employment. In particular, employment in construction occupations is affected quickly by a supply shock. Our median responses imply that one quarter after a positive supply shock, 80 percent of the new jobs in construction and production have already been created - compared to 50 percent of all new jobs. Later in the build-up of a boom or bust, employment in other occupations then changes in a more pronounced way. Thus, supply shocks are able to explain why occupational employment ratios are leading aggregate employment in the data.

– insert Figure 6 here –

Age. Our estimates suggest that both shocks contribute to the relatively strongly pro-cyclical employment of young workers. Supply shocks appear more important quantitatively. We see this in the left column of Figure 6 (a) and (b), where we compare above 25 year olds to 16-24 year olds.

After positive supply shocks, the employment of older employees falls significantly relative to the employment of 16-24 year olds, so young workers are hired overproportionately. The responses to demand shocks go into the same direction but are less pronounced quantitatively.

Our median responses imply that, at the peak of aggregate employment after a supply shock, about 36 percent of new full-time employment is created among the 16-24 year old. As a comparison, this age group amounts to roughly 10 percent of full-time employment in the long-run mean. At the aggregate peak response after a demand shock, about 22 percent of the new full-time employment is created among the 16-24 year old.

Considering the transmission of shocks, we observe that our identified shocks can not explain much of the lead of the employment ratios built by age compared to aggregate employment. In fact, many estimated responses in Figure 6 appear rather synchronized to the responses of aggregate employment. There are some strong median reactions on impact but they are statistically insignificant. The strongest reactions which are statistically significant occur three to four quarters after the shock, so at about the time where also aggregate employment peaks.

– insert Figure 7 here –

Race/ethnic origin. We now consider the composition of employment by race/ethnic origin, displayed in Figure 7. Here, we find that fluctuations occur primarily in employment ratios between the different groups of non-whites on the one hand and whites on the other hand. Further, demand shocks are found to be the main source of such fluctuations. After positive demand shocks, the employment of blacks and Asians rises significantly relative to the employment of whites. The response of Hispanics relative to whites to demand shocks is similar, though only marginally significant. Our median responses imply that, at the peak response of aggregate employment after a demand shock, about 42 percent of the newly created full-time jobs are taken up by non-whites who are responsible for about 28 percent of all full-time employment in the long-run mean. After supply shocks, we find a significant rise in the employment of Hispanics relative to whites. For the other ratios, we only find very short-lived significant responses to supply shocks.

Concerning the timing of the statistically significant reactions to demand shocks, we find that they are strongest after about four to five quarters. Thus, in this dimension of disaggregation, the

employment ratios are even slightly lagging behind aggregate employment after demand shocks. Supply shocks can explain some of the lead structure we find in the unconditional moments but are else relatively unimportant in this dimension of disaggregation.

– insert Figure 8 here –

Education. Finally, we consider groups built by educational attainment. Here, we observe that especially college graduates have a less strongly pro-cyclical employment. Our results suggest that this is primarily driven by supply shocks. For example, employment of high-school graduates rises significantly relative to that of college graduates after a supply shock while demand shocks only cause a very late, weak, and less significant positive response that stands in contrast to the unconditional correlation, see Figure 8 (a) and (b), lower panels. We make similar observations when we compare individuals with some college education to college graduates. Quantitatively, our median responses imply that, at the aggregate peak after a supply shock, about 86 percent of the newly created full-time jobs go to people with less than a college-degree, compared to a long-run share of people with this education in total full-time employment of about 65 percent.

Concerning timing, we find the strongest significant responses to supply shocks three quarters after the shock, thus one quarter before the peak reaction in aggregate employment. Consequently, supply shocks can explain a small part of the leads of the employment ratios built by education relative to aggregate employment in the unconditional moments.

Summary. Considering the summary of our estimation results in Table 3, we see that the employment of males, young people, non-whites, the less educated, and workers in blue-collar occupations reacts more strongly to business-cycle shocks. Another general insight is that supply shocks are more important to explain the fluctuations of the employment composition. Especially employment ratios built by gender, occupation, age, and education react more strongly and more significantly to supply shocks than to demand shocks. By contrast, demand shocks are rather important to explain fluctuations in employment ratios built by employment status and by race/ethnic origin.

A further interesting general insight from our results concerns the timing of the responses to supply and demand shocks. Often, we find that supply shocks cause pronounced and significant effects on the employment ratios earlier than demand shocks do, see the periods with significant

reactions and those with the strongest reactions in Table 3. For example, there is not a single significant impact response to demand shocks while there are several to supply shocks. Thus, the fact that the employment composition generally leads aggregate employment in the data can primarily be attributed to supply shocks.

Implications. We highlight three implications from these results. First, our results have interesting distributional implications. The groups which we have found to be hit hardest in recessions are relatively poorer population groups. This is in particular true for supply-side driven recessions. Thus, poorer groups bear the lion's share of the costs of the business cycle. Further, inequality tends to rise in such recessions. This, in turn, increases the social costs of recessions.

Second, these distributional consequences induce the question whether policy can stabilize the composition of employment over the business cycle. To approach this question, let us perform a thought experiment where the government uses demand-side impulses (such as increases in government spending) to counter negative supply shocks. Though we are aware of the limits of interpreting shocks as a policy reaction which might - to a certain degree - be systematic, we find this thought experiment insightful to explore what demand-side policies can or can not achieve.

Our results imply that demand-side innovations can cushion the effects of supply shocks only to a limited degree. Especially, demand-side policies do not seem adequate to offset the particularly strong job losses in blue-collar occupations in a supply-side driven recession. The same holds true for the job losses suffered by the less educated. In other cases such as gender or age, demand shocks can weaken, though not completely offset, the changes in the employment composition caused by supply shocks. Further, the different timing of the effects of supply and demand shocks implies that demand-side policies cannot impede the early lay-offs caused by negative supply shocks. As, for many ratios, the strongest effects of demand shocks come with a substantial delay, demand-side reactions would have to actually lead the supply shocks they try to accommodate.

Third, our findings help to improve our understanding of the chain of events in the transmission of business-cycle shocks, especially of those that originate on the supply side. The transmission of supply shocks seems to operate rather early through sectors with high shares of male, blue-collar employment. In such sectors, people seem to lose their jobs rather early in a recession, before it spreads over to other parts of the economy.

Potential explanations. Hoynes et al. (2012) relate the different employment volatilities to the demographic composition of sectoral employments and the different importance of hiring rates for different age groups. Similar arguments can be applied to think about what can lie behind our novel findings concerning the timing and sources of heterogeneous employment fluctuations.

The dynamics in employment ratios between genders and education groups can be related to the higher likelihood of men and the less educated to work in construction-intense sectors. We document that employment dynamics in construction occupations are mainly driven by supply shocks. It appears likely that the same is true for the construction sector. The sectoral composition of employment among men and the less educated can then also shed some light why supply shocks are important for the dynamics of employment ratios built by gender and education. We can also attribute the timing of the reaction of these employment ratios to the sectoral composition of employment. Considering the correlation function with aggregate employment, we find that activity in the construction sector is leading the aggregate state of the cycle.⁵ This can thus give a hint to understand the early reactions of men and the less educated.

To think about reasons for the dynamics in employment of different races, it is helpful to bear in mind that the public sector is an important employer for blacks. Demand shocks include changes in government spending, a substantial part of which is public employment. The high share of blacks in public employment can thus be one reason why we find that demand shocks affect foremost black workers. Further, changes in public employment lag behind aggregate employment.⁶ This can explain why we do find delayed responses of employment ratios between races to demand shocks.

Finally, to understand the dynamics of employment ratios between age groups, we argue that the young, more than other groups, are affected by changes in hiring rates. Hiring rates in the total economy are leading the aggregate state of the cycle.⁷ This can help to understand why employment ratios built by age also lead the cycle.

⁵As measures of activity in the construction sector, we consider total private residential investment (Source: BEA, Series ID: PRFI) and total construction (Source: OECD, Series ID: PRCNTO01USQ661S). Series are divided by Civilian Noninstitutional Population, logged and de-trended (HP 1600). We find that they lead the cycle by six and two quarters, respectively.

⁶Using the series 'All Employees: Government' (Source: BLS, Series ID: USGOVT, divided by the Civilian Noninstitutional Population, logged and de-trended (HP 1600)), we find that it has the strongest correlation with aggregate employment at lag $i = 5$.

⁷We find that total private hires (Source: BLS, Series ID: JTS1000HIL, divided by Civilian Noninstitutional Population, logged and de-trended (HP 1600)) lead aggregate employment by two quarters.

5 Conclusion

Whose jobs are cut in a recession and for whom are jobs created in booms? In this paper, we find that the employment of males, young people, non-whites, the less educated, and workers in blue-collar occupations fluctuates more heavily over the business cycle. Thus, mostly jobs of these population groups are cut in recessions and created in booms.

Who is fired first in a recession and who finds a job early in a boom? In this paper, we find that groups whose employment fluctuates more heavily over the cycle also tend to be affected early in the transmission process. In recessions, particularly men, the less educated and workers in blue-collar occupations are laid off early before the recession spreads to other population groups. Symmetrically, these are the groups which find jobs early in booms.

What are the drivers of heterogeneous employment dynamics across population groups? In this paper, we disentangle between the effects of supply and demand shocks. Supply shocks primarily drive differences between occupations, genders, age groups, and education groups. Demand shock are important for the dynamics of employment ratios between races and ethnic origins. The fact that groups with high employment fluctuations tend to be affected at the beginning of the business cycle is found to be foremost driven by supply-side disturbances.

References

- Abell, John D., 1991. Distributional Effects of Monetary and Fiscal Policy: Impacts on Unemployment Rates Disaggregated By Race and Gender. *American Journal of Economics and Sociology* 50, 269-284.
- Carpenter, Seth B., and William M. Rodgers, 2004. The disparate labor market impacts of monetary policy. *Journal of Policy Analysis and Management* 23, 813-830.
- Clark, Kim B., and Lawrence H. Summers, 1981. Demographic Differences in Cyclical Employment Variation. *Journal of Human Resources* 16, 61-79.
- Doan, Thomas, Robert Litterman, and Christopher Sims, 1984. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews* 3, 1-100.
- Elsby, Michael W., Bart Hobijn, and Aysegul Sahin, 2010. The Labor Market in the Great Recession. *Brookings Papers on Economic Activity* 41, 1-43.
- Giavazzi, Francesco, and Michael McMahon, 2013. The Household Effects of Government Spending. Chapter in *Fiscal Policy after the Financial Crisis*, edited by Alberto Alesina and Francesco Giavazzi, 103 - 141.
- Hoynes, Hilary W., Douglas L. Miller, and Jessamyn Schaller, 2012. Who Suffers during Recessions? *Journal of Economic Perspectives* 26, 27-48.

- Hines Jr, James R., Hilary W. Hoynes, and Alan B. Krueger, 2001. Another Look at Whether a Rising Tide Lifts All Boats. NBER Working Papers 8412.
- Jefferson, Philip N., 2005. Does monetary policy affect relative educational unemployment rates. *American Economic Review* 95, 76-82.
- Jefferson, Philip N., 2008. Educational Attainment and the Cyclical Sensitivity of Employment. *Journal of Business and Economic Statistics* 26, 526-535.
- Rodgers, William M., 2008. African American and white differences in the impacts of monetary policy on the duration of unemployment. *American Economic Review* 98, 382-386.
- Thorbecke, Willem, 2001. Estimating the effects of disinflationary monetary policy on minorities. *Journal of Policy Modeling* 23, 51-66.
- Uhlig, Harald, 2005. What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics* 52, 381-419.
- Williams, Roger C., 2004. Monetary policy and unemployment: A disaggregated analysis. *International Advances in Economic Research* 10, 180-190.
- Zavodny, Madeline, and Tao Zha, 2000. Monetary policy and racial unemployment rates. *Economic Review*, Federal Reserve Bank of Atlanta, Q4, 1-16.

Tables and Figures

Table 1: Unconditional correlations of group employment ratios to aggregate employment

		contemp.	strongest	at lag
0	Aggregate Employment	1.00	1.00	0
1	Women vs. Men	-0.45	-0.87	-3
2	Part time vs. Full time	-0.73	-0.92	-2
3	Services vs. Management, Professionals	-0.02	-0.38	8
4	Sales and Office vs. Management, Professionals	0.17	0.56	-5
5	Sales and Office vs. Services	0.15	0.19	-2
6	Construction vs. Management, Professionals	0.52	0.83	-5
7	Construction vs. Services	0.54	0.73	-3
8	Construction vs. Sales and Office	0.57	0.79	-3
9	Production vs. Management, Professionals	0.37	0.82	-3
10	Production vs. Services	0.38	0.70	-3
11	Production vs. Sales and Office	0.34	0.71	-3
12	Production vs. Construction	-0.23	-0.32	-7
13	Age 25-34 vs. Age 16-24	-0.71	-0.82	-2
14	Age 35-44 vs. Age 16-24	-0.69	-0.86	-2
15	Age 35-44 vs. Age 25-34	0.19	0.29	1
16	Age 45-54 vs. Age 16-24	-0.73	-0.89	-2
17	Age 45-54 vs. Age 25-34	-0.37	-0.55	-2
18	Age 45-54 vs. Age 35-44	-0.57	-0.63	-2
19	Age 55+ vs. Age 16-24	-0.71	-0.88	-2
20	Age 55+ vs. Age 25-34	-0.30	-0.53	-3
21	Age 55+ vs. Age 35-44	-0.44	-0.54	-2
22	Age 55+ vs. Age 45-54	-0.06	-0.26	-4
23	Black vs. White	0.65	0.73	-1
24	Asian vs. White	0.47	0.48	1
25	Asian vs. Black	0.16	-0.36	-8
26	Hispanic vs. White	0.50	0.68	-3
27	Hispanic vs. Black	-0.03	0.36	-6
28	Hispanic vs. Asian	-0.16	0.45	-8
29	High School graduate vs. High School Drop-out	-0.14	-0.45	-7
30	Some College vs. High School Drop-out	-0.08	-0.57	-7
31	Some College vs. High School Graduate	0.11	-0.43	-6
32	College Graduate vs. High School Drop-out	-0.22	-0.64	-6
33	College Graduate vs. High School Graduate	-0.18	-0.51	-6
34	College Graduate vs. Some College	-0.29	0.35	8

Notes: The table displays results of the cross-correlation function between the respective employment ratio $EmpRatio$ and aggregate employment $AggEmp$, that is $Corr(EmpRatio_t, AggEmp_{t-i})$. Column 3 displays the contemporaneous correlation, $Corr(EmpRatio_t, AggEmp_t)$. Columns 4 and 5 display the value of the strongest correlation within a ± 8 quarters window together with the lag i at which the strongest correlation is found.

Table 2: Imposed sign restrictions.

	Reaction in quarters $t = 0 - 3$ after the shock				
	labor productivity	GDP	aggregate employment	real interest rate	employment ratio
productivity shock	> 0	> 0			
non-productivity shock	≤ 0	≥ 0	≥ 0		

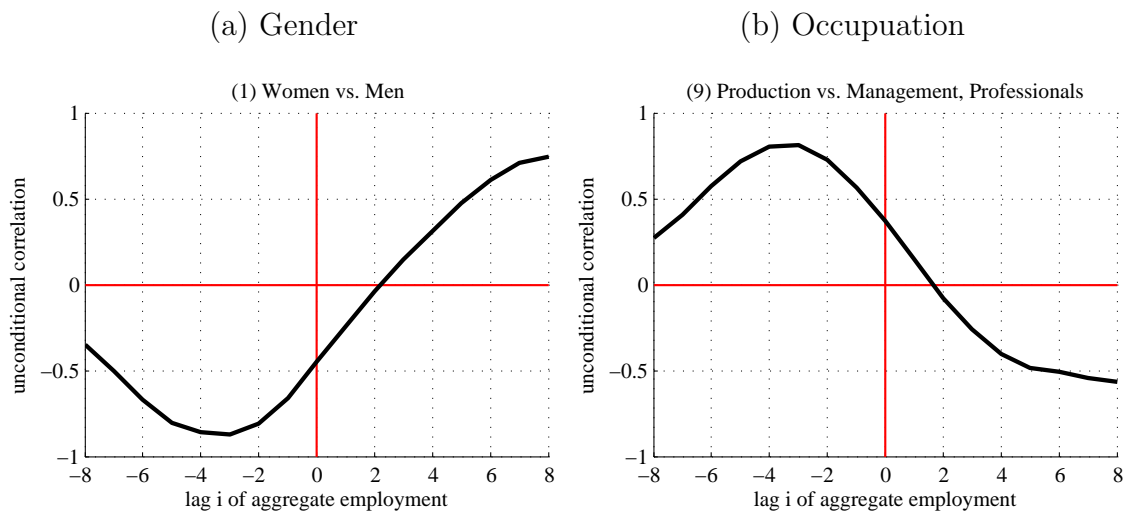
Notes: The table refers to positive shocks. For negative shocks, all signs are reversed. No entry means no sign restriction on this variable. Periods t measured in quarters after shock, where $t = 0$ is the impact period.

Table 3: Summary of significant estimation results.

		Supply shock				Demand shock			
		signif.	periods	peak	at $t =$	signif.	periods	peak	at $t =$
0	Aggregate Employment	**	2-5	1.00	4	***	0-6	1.00	3
1	Women vs. Men	***	1-4	-1.09	1	**	2-4	-0.51	3
2	Part time vs. Full time	**	3-6	-2.30	4	***	2-7	-2.18	4
6	Construction vs. Management, Prof.	***	0-10	4.76	0	-	-	-	-
7	Construction vs. Services	***	0-8	7.08	0	-	-	-	-
8	Construction vs. Sales, Office	***	0-9	3.41	0	-	-	-	-
9	Production vs. Management, Prof.	**	3-4	2.84	3	**	2-4	2.06	2
10	Production vs. Services	-	-	-	-	*	4-4	0.81	4
11	Production vs. Sales, Office	**	3-5	1.92	3	*	4-4	0.66	4
12	Production vs. Construction	*	0-0	-4.28	0	-	-	-	-
13	Age 25-34 vs. Age 16-24	**	3-8	-2.62	5	*	3-4	-1.21	4
14	Age 35-44 vs. Age 16-24	***	2-8	-2.41	4	**	1-5	-1.32	3
15	Age 35-44 vs. Age 25-34	*	7-8	0.38	7	-	-	-	-
16	Age 45-54 vs. Age 16-24	***	1-8	-3.02	4	**	2-5	-1.48	4
18	Age 45-54 vs. Age 35-44	**	4-6	-0.71	4	**	4-7	-0.37	4
19	Age 55+ vs. Age 16-24	***	0-8	-4.52	0	**	1-5	-1.47	4
20	Age 55+ vs. Age 25-34	*	2-4	-1.87	2	-	-	-	-
21	Age 55+ vs. Age 35-44	**	2-5	-1.46	3	*	4-5	-0.52	4
23	Black vs. White	-	-	-	-	***	3-7	0.98	5
24	Asian vs. White	**	0-0	-6.43	0	***	2-9	1.02	4
25	Asian vs. Black	**	0-0	-6.83	0	**	4-9	0.41	4
26	Hispanic vs. White	**	0-7	3.76	0	*	4-5	0.65	4
27	Hispanic vs. Black	**	0-0	3.81	0	-	-	-	-
28	Hispanic vs. Asian	***	0-1	7.44	0	*	5-9	-0.51	6
32	College Graduate vs. HS Drop-Out	**	2-7	-1.76	3	-	-	-	-
33	College Graduate vs. HS Graduate	***	3-7	-1.64	3	**	10-13	0.38	10
34	College Graduate vs. Some College	***	3-7	-1.08	3	*	11-13	0.25	11

Notes: 'signif.' gives the significance level, '*', '**', or '***' indicate that the 68%, 80%, or 90% confidence bands include zero. 'periods' lists the quarters after the shock with significant (68%) responses, where $t = 0$ is the impact period. 'peak' gives the strongest significant (68%) median response. 'at $t =$ ' gives the period of the 'peak' response in quarters after the shock. Employment ratios not listed in the table do not show significant responses to neither shock. 'Prof.' is the abbreviation for Professionals, 'HS' stands for High School.


Figure 1: Unconditional correlation functions of employment ratios with aggregate employment.




Notes: The solid lines show the correlation function between the respective employment ratio $EmpRatio$ and aggregate employment $AggEmp$, that is $Corr(EmpRatio_t, AggEmp_{t-i})$, within a ± 8 quarters window.

Figure 2: Impulse responses of labor productivity, GDP, employment, and the real interest rate.

(a) Supply shock




(b) Demand shock



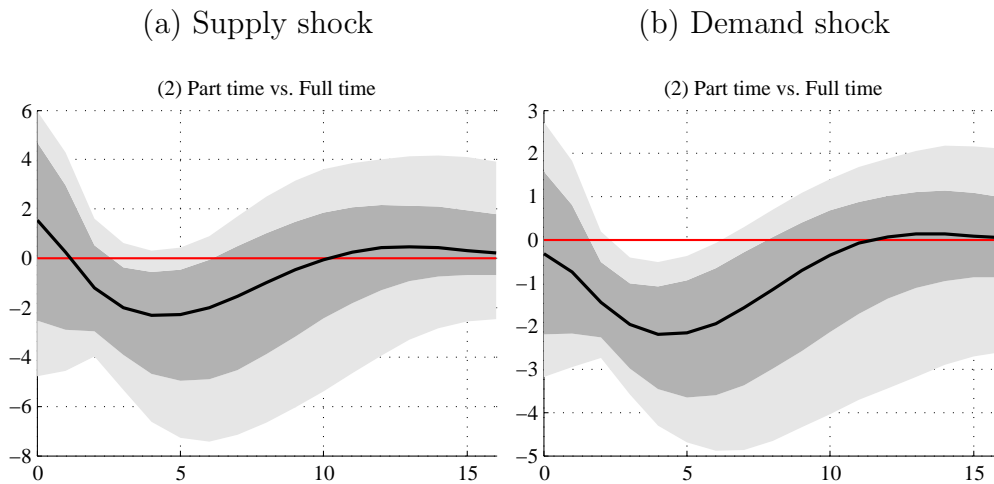
Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Figure 3: Responses of the female-male employment ratio.



Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.


Figure 4: Responses of the part-time/full-time employment ratio.



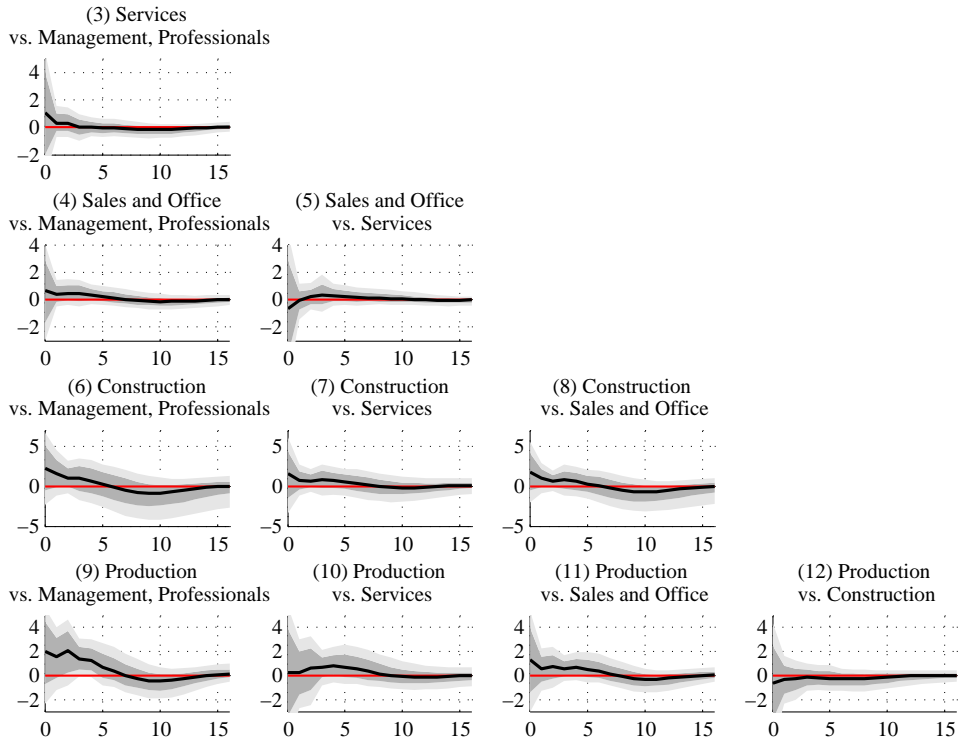
Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Figure 5: Responses of employment ratios built by occupation.

(a) Supply shock




(b) Demand shock



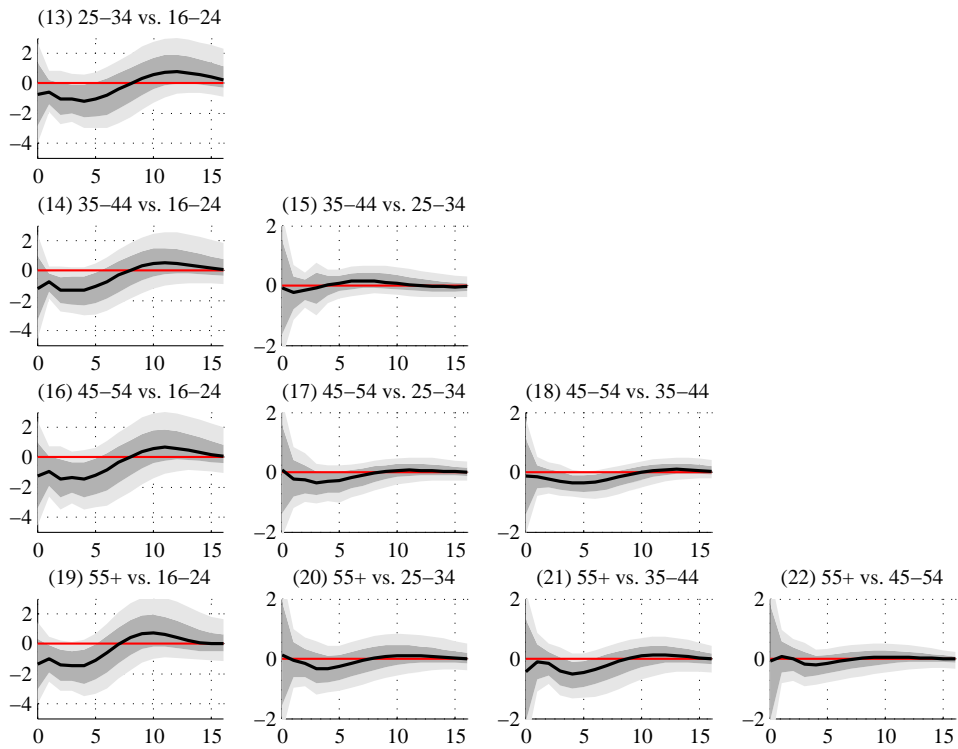
Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Figure 6: Responses of employment ratios built by age.

(a) Supply shock



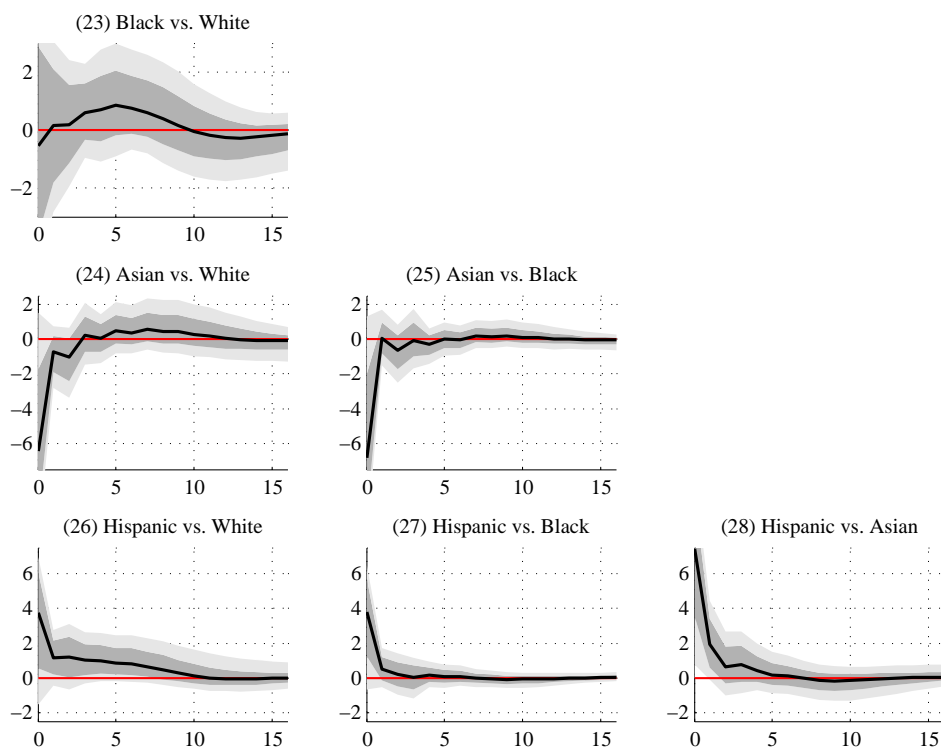
(b) Demand shock




Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Figure 7: Responses of employment ratios built by race/ethnic origin.

(a) Supply shock



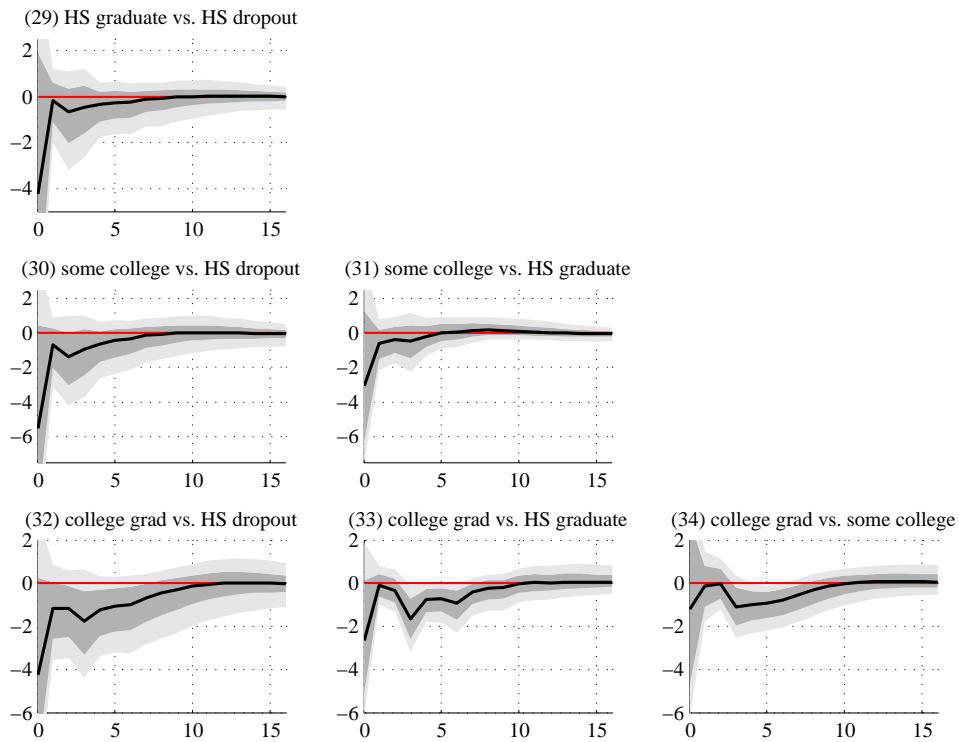
(b) Demand shock



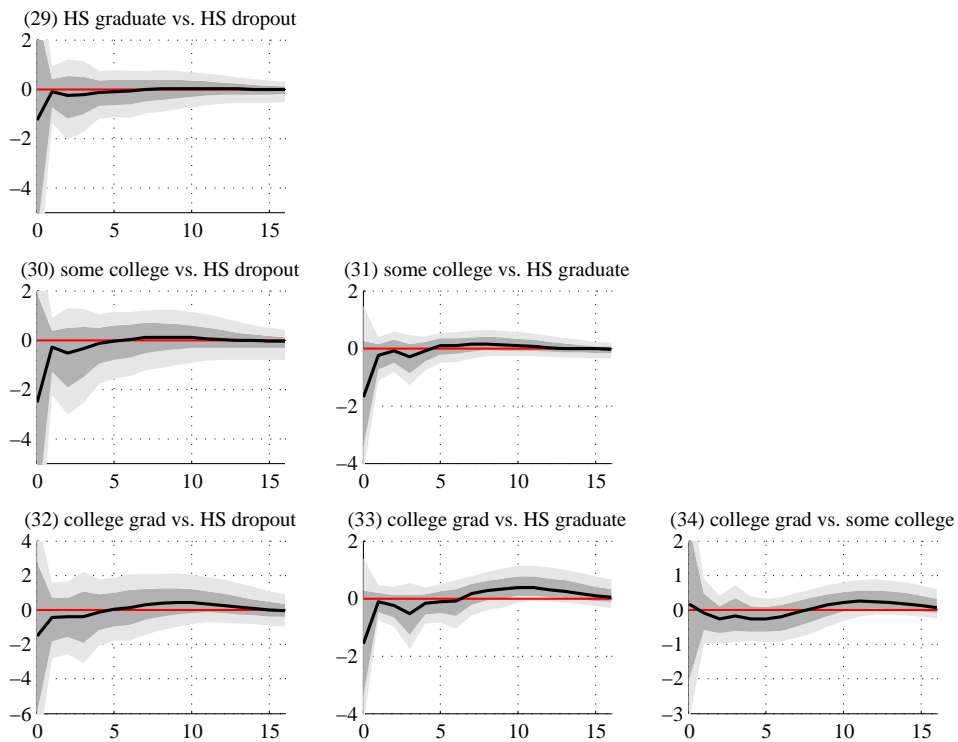
Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Figure 8: Responses of employment ratios built by educational attainment.

(a) Supply shock



(b) Demand shock



Notes: Solid lines represent the median responses and the grey areas are the 68% (light grey) and 90% (dark grey) confidence bands, respectively. Numbers on the horizontal axes denote quarters after the shock and the responses are expressed in percentage points.

Appendix

Table A: Data sources: Employment

Dimension	Group	Series Title	Series ID
-	Aggregate	Civilian Employment-Population Ratio	LNS12300000
Gender	Men	(Seas) Employment-Population Ratio - Men	LNS12300001
	Women	(Seas) Employment-Population Ratio - Women	LNS12300002
Employment	Full time	(Seas) Employed, Usually Work Full Time	LNS12500000Q
Status	Part time	(Seas) Employed, Usually Work Part Time	LNS12600000Q
Occupation	Management, Professionals	(unadj)- Employed full time, Wage and salary workers, Management, professional, and related occupations	LEU0254471200
	Services	(unadj)- Employed full time, Wage and salary workers, Service occupations	LEU0254490000
	Sales and Office	(unadj)- Employed full time, Wage and salary workers, Sales and office occupations	LEU0254496800
	Construction	(unadj)- Employed full time, Wage and salary workers, Natural resources, construction, and maintenance occupations	LEU0254504000
	Production	(unadj)- Employed full time, Wage and salary workers, Production, transportation, and material moving occupations	LEU0254512800
Age	16-24	(unadj)- Employed full time, Wage and salary workers, 16 to 24 years	LEU0252886200
	25-34	(unadj)- Employed full time, Wage and salary workers, 25 to 34 years	LEU0252888400
	35-44	(unadj)- Employed full time, Wage and salary workers, 35 to 44 years	LEU0252889000
	45-54	(unadj)- Employed full time, Wage and salary workers, 45 to 54 years	LEU0252889600
	55+	(unadj)- Employed full time, Wage and salary workers, 55 years and over	LEU0252890200
Race/Ethnic Origin	White	(unadj)- Employed full time, Wage and salary workers, White	LEU0254898500
	Black	(unadj)- Employed full time, Wage and salary workers, Black or African American	LEU0254898800
	Asian	(unadj)- Employed full time, Wage and salary workers, Asian	LEU0254468100
	Hispanic	(unadj)- Employed full time, Wage and salary workers, Hispanic or Latino	LEU0254899100
Education	HS dropout	(unadj)- Employed full time, Wage and salary workers, Less than a high school diploma, 25 years and over	LEU0252916400
	HS graduate	(unadj)- Employed full time, Wage and salary workers, High school graduates, no college, 25 years and over	LEU0252917000
	some college	(unadj)- Employed full time, Wage and salary workers, Some college or associate degree, 25 years and over	LEU0254929100
	college grad	(unadj)- Employed full time, Wage and salary workers,	LEU0252918200

Notes: Series marked with (Seas) is provided in seasonally adjusted form by the BLS. We seasonally adjust the series marked with '(unadj-)' using X-12-ARIMA. All series are from the U.S. Bureau of Labor Statistics (BLS).

Table B: Data sources: Others

Series Title	Series ID	Source
Gross Domestic Product	GDP	BEA
Gross Domestic Product: Implicit Price Deflator	GDPDEF	BEA
Civilian Noninstitutional Population	CNP16OV	BLS
Nonfarm Business Sector: Output Per Hour of All Persons	OPHNFB	BLS
Effective Federal Funds Rate	FEDFUNDS	BFED

Notes: BEA: U.S. Bureau of Economic Analysis, BLS: U.S. Bureau of Labor Statistics, BFED: Board of Governors of the Federal Reserve System.

Table C: Definition of data variables

Time series	Construction	Description
Output	$\text{hp} \left(\log \left(\frac{(GDP)}{(GDPDEF)(CNP16OV)} \right) \right)$	HP filtered cyclical component of log of real GDP per capita
Labor productivity	$\text{hp} (\log (OPHNFB))$	HP filtered cyclical component of log labor productivity
Real interest rate	$\text{hp} \left(\frac{FEDFUNDS}{100} - \log \left(\frac{GDPDEF(+1)}{GDDEF} \right) * 4 \right)$	HP filtered cyclical component of annualized real interest rate
Employment	$\text{hp} (\log (LNS12300000))$	HP filtered cyclical component of aggregate employment
Employment ratio	$\text{hp} \left(\log \left(\frac{Emp_i}{Emp_j} \right) \right)$	HP filtered cyclical component of the employment ratio between group i versus j

Notes: The function hp stands for Hodrick-Prescott (HP) filter and computes the cyclical component of a time series using the HP filter with a smoothing parameter of 1600.

