The Long-Lived Cyclicality of the Labor Force Participation Rate*

Tomaz Cajner John Coglianese Joshua Montes

November 25, 2019

Preliminary; conference draft. Comments welcome.

Abstract

How cyclical is the labor force participation rate (LFPR)? We measure cyclicality by examining the response of state-level LFPRs to exogenous local business cycle shocks, which sidesteps several challenges with measuring LFPR cyclicality from aggregate time series. We find that the LFPR is highly cyclical, but with a significantly longer lag structure than the unemployment rate. After a negative shock, the LFPR steadily declines and does not reach its trough for 4–5 years, well after the unemployment rate has begun to recover. The LFPR remains significantly depressed in the long run after a shock, although this merely reflects changes in the composition of the population induced by the shock, as the demographically-adjusted LFPR fully recovers. We show that shocks spur relative declines in the population of 25–39 year olds in affected states, likely due to migration, which skews the composition of the population towards groups with lower LFPRs. Further, the response of LFPR varies considerably across skill groups, showing larger cyclicality among men, younger workers, and workers with less education. We conclude that failing to account for the cyclical dynamics of LFPR and focusing only on the unemployment rate is likely to give an incomplete picture of the labor market response to business cycles.

Keywords: labor force participation, labor supply, labor force composition, labor force demographics, full employment, Okun's law, geographic mobility, labor mobility, regional migration

JEL Classification: E24, J21, J22, J61, J64

*All authors are at the Federal Reserve Board of Governors. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

1 Introduction

How does the US labor market respond to business cycle shocks? One perspective is to focus on the unemployment rate and ignore other labor market outcomes. However, many observers have noted that the labor force participation rate (LFPR)—the share of adults either working or looking for work—exhibits some degree of cyclicality as well (see Montes (2018) and references cited therein). Measuring the degree of cyclicality in the LFPR is complicated, though, by the presence of trend movements unrelated to the business cycle, including the entry of women into the workforce in past decades and the aging of the baby boom generation. Since many observers disagree about the exact magnitude of these trends, there is also substantial disagreement about the extent of cyclicality in labor force participation.

We examine the cyclicality of labor force participation through a different lens—measuring its response to state-level business cycles. In doing so, we sidestep the issue of identifying trend changes in labor force participation. Our research design identifies the response of labor market outcomes to unexpected declines in state output without imposing strict parametric assumptions or assuming that cyclical effects dissipate in the long run.

We find that labor force participation *is* cyclical, but with a different pattern than unemployment. Compared to the response of the unemployment rate to a business cycle shock, the LFPR responds with a smaller elasticity, a more delayed impact, and a more persistent effect. The LFPR remains below its pre-shock level even 10–15 years after the initial shock, even after the unemployment rate fully recovers. This long-run decline is driven by a change in the composition of the population, though. Shocks lead to increased out-migration for adults ages 25–39, who tend to have higher participation rates than average, which mechanically pulls down the participation rate in the long-run. The size of labor force participation and unemployment responses, along with their timing, vary substantially across demographic groups. These facts together suggest that studying the aggregate unemployment rate on its own without accounting for labor force participation and its diverse changes across groups offers an incomplete picture of the labor market response to business cycle shocks.

Our approach uses state-level variation in business cycles to estimate the cyclicality of labor force participation. We measure cyclical variation at the state level using Gross State Product (GSP) and estimate the response of labor market outcomes using the local projections method (Jordà, 2005). To avoid endogeneity between output and the labor market, we instrument for changes in GSP with a shift-share instrument exploiting variation in local exposure to national changes in output across industries (Bartik, 1991). We verify that this instrument is not a weak instrument and is not subject to concerns about inference with shift-share instruments (Adão, Kolesár, and Morales, 2018).

We start by estimating the response of the unemployment rate. When a state is hit by a temporary, contractionary 1 percent shock to output, the unemployment rate rises 0.3 p.p. immediately and peaks at about 0.45 p.p. above its pre-shock level 1–2 years after the shock. This estimate is squarely in the range of Okun's law coefficients estimated in the literature, confirming that our method measures cyclicality accurately (Ball, Leigh, and Loungani, 2017). After peaking, the unemployment rate gradually declines until it returns to its pre-shock level, fully recovering about 6 years after the shock.

The LFPR also responds to cyclical shocks, but with a substantial delay compared to the unemployment rate. The LFPR declines by only a small amount initially, but eventually falls 0.2–0.3 p.p. below its pre-shock level by 4–6 years after the shock. The delayed response of the LFPR means that there is a multi-year period in which the LFPR is continuing to decline even as the unemployment rate is recovering. Only after the unemployment rate is fully recovered does the participation rate begin to recover. The delay between recovery in these two outcomes means that observers who focus only on the unemployment rate may underestimate the extent of slack remaining in the labor market after a recession, particularly in the period approximately 6 years or more after the initial shock.

In addition to delayed cyclicality, the state-level LFPR fails to fully recover from shocks in the long-run. Even 10 years after the shock, the LFPR remains well below its pre-shock level. Accordingly, EPOP remains below its pre-shock level in the long-run too. The lack of full recovery is surprising given the full recovery of unemployment, and it admits the possibility that cyclical shocks may have permanent effects on employment even if they do not lead to permanent changes in unemployment.

What could be responsible for the persistent shortfall of the LFPR in shocked states? There are several different possibilities. It could represent hysteresis effects in the labor market, as individuals who are laid off during a cyclical downturn become permanently cut off from work and employment falls permanently. Similarly, it could represent individuals finding productive activities outside of the labor force, including retiring early, going back to school, or staying at home for child care. Both of these forces would lower employment rates conditional on the demographic makeup of the population, but an alternative possibility is that the demographic makeup itself could change in response to the shock. Shocks could induce high-participation groups to move away and shift the population towards low-participation groups without permanently affecting the participation rates within any group.

We show that the persistent decline in the state LFPR is driven by changes in the demographic makeup of the population. Using individual-level data, we separate the participation rate at the state level into a demographic component explained by observable characteristics of the individuals in that state and a demographically-adjusted component containing the residual LFPR. The persistent decline is driven entirely by the demographic component, indicating that the population has changed in favor of low-participation groups. Decomposing the demographic component reveals that the primary changes are with respect to age and gender, and to a lesser extent education.

Accordingly, the demographically-adjusted LFPR fully recovers from shocks, suggesting limited scope for hysteresis at the state level. The demographically-adjusted component of the LFPR follows a similar pattern as the overall LFPR initially, declining with a delay compared to the unemployment rate and reaching its trough about 4 years after the shock. After reaching this level, though, the demographically-adjusted LFPR begins to steadily rise again and fully recovers by 8 years after the shock. However, this full recovery still comes with a delay compared to the unemployment rate.

The response of the prime-age LFPR confirms the delayed full recovery of the demographically-adjusted LFPR. We repeat our analysis focusing only on the LFPR for individuals ages 25–54. This prime-age LFPR falls in response to the shock by a similar magnitude as the overall LFPR and reaches its trough with a similarly long delay. However, after reaching its trough, the prime-age LFPR begins to rise and eventually fully recovers by about 8 years after the shock. Combined with the evidence above, this indicates that the persistent decline in the overall LFPR we estimate can be attributed to changes in the relative prevalence of prime-age individuals in the population as a result of the shock, rather than permanent changes in participation or composition within the group of prime-age individuals.

We provide evidence in favor of one possible explanation for compositional changes—out-migration of high-participation individuals. We examine which age groups are responsible for the changes in composition using data on state-level population by single-year age group. Our estimates indicate that business cycle shocks lead to declines in population of 25–39 year olds in the states affected by the shocks, while the rest of the population is largely unaffected. Workers in this age range tend to have substantially higher participation rates on average, so declines in this part of the population will naturally tend to pull down the LFPR mechanically.

We also document heterogeneity in the response of labor market outcomes across different demographic groups. For prime-age workers, we find that men see a larger initial increase in the unemployment rate than women, but the recovery of unemployment afterwards is much faster for men. Younger workers (ages 16–24) exhibit much larger cyclical responses of unemployment and participation than the overall population, while older workers (ages 55+) see lower cyclicality of both. Our estimates also show a sharp difference by education level with less-educated prime-age workers experiencing a large increase in unemployment and decrease in participation after a shock, while more-educated workers experience only a small, short-lived tick up in unemployment and no significant change in participation.

Our paper is related to several strands of literature. First, several recent papers provide estimates of LFPR cyclicality (Aaronson, Cajner, Fallick, Galbis-Reig, Smith, and Wascher, 2014; Council of Economic Advisers, 2014; Erceg and Levin, 2014; Balakrishnan, Dao, Sole, and Zook, 2015; Montes, 2018). While these papers do argue that the cyclical response of LFPR can be delayed, one of the main contributions of our paper is to use a method that is particularly well-suited for estimating long lags in LFPR cyclicality. More precisely, unlike the previous papers in this literature we estimate the dynamic response of LFPR to output shocks by using the local projection regressions, which allow for the possibility of very persistent effects on LFPR. Moreover, by using a shift-share instrumental variable approach, we are able to establish a causal link between business cycle shocks and the dynamic response of LFPR (Balakrishnan, Dao, Sole, and Zook (2015) also use a similar shift-share instruments, but for employment, while our paper uses output). Second, following the early work of Blanchard and Katz (1992), several papers investigate adjustments to economic shocks at the state level (Dao, Furceri, and Loungani, 2017; Amior and Manning, 2018; Yagan, 2019). Our contribution to this literature is to provide estimates of the importance of compositional effects and to show how migration patterns differ by age. Third, we built on the recent methodological contributions in the shift-share empirical design (Adão, Kolesár, and Morales, 2018; Borusyak, Hull, and Jaravel, 2018; Goldsmith-Pinkham, Sorkin, and Swift, 2018).

2 Research Design

We measure the cyclicality of labor force participation by estimating its response to business cycle shocks. Instead of conducting our analysis at the national level, we examine state-level business cycles to sidestep the issue of trend changes in participation. In this section, we outline our research design, starting with the identification problem and our approach to solving it. We then turn to the issue of inference and the description of the data we use in this analysis.

2.1 Identification

Estimating the dynamic cyclical responses of national outcomes typically requires strict assumptions. For example, time series models usually need a mean zero cyclical component, which rules out hysteresis by definition. Further, identification in those models relies on a trend component that is smooth and identifiable—a strong assumption for LFPR, given the sharp and changing nature of those trends for various subgroups of the population.

To meet these challenges, we use state-level panel data to estimate the dynamic cyclical responses of labor market outcomes to a state-level business cycle shock using the local projections method. In particular, we measure the impulse response functions (IRFs) of a shock by estimating the following series of k regressions:

$$y_{s,t+k} - y_{s,t-1} = \beta^{(k)}\text{Shock}_{s,t} + \Theta X_{s,t} + \epsilon_{s,t+k} \tag{1}$$

where $y_{s,t}$ represents the labor market dependent variable of interest—the unemployment rate, the LPFR, or the EPOP—of state s in time t; k indexes the regression that measures the effect of the shock at time t on the dependent variable $t+k$ periods ahead; $\text{Shock}_{s,t}$ is the measure of the business cycle shock (defined below); and $X_{s,t}$—which does not vary by group j—represents a vector of control variables. In our baseline specification, the controls include state fixed effects, year fixed effects, and the growth rate of the state-level home price index.

Our local projections regressions control for national trends through the inclusion of year fixed effects. This does not impose strict assumptions about the smoothness of trends, as would be needed in national-level time series regressions. Any nation-wide trends that affect labor market outcomes across all states equally, including demographic shifts such as the aging of the baby boom generation, are controlled for nonparametrically by this approach.

Shocks: We measure the business cycle shock using gross state-level product (GSP). Specifically, we define $\text{Shock}_{s,t} \equiv \Delta GSP_{s,t}$, where $\Delta GSP_{s,t}$ is the year-over-year percentage change in real gross state product. GSP measures the total output of goods and services produced in each state; changes in GSP provide a natural measure of business cycle fluctuations at the state level.

Defining business cycles based on output is more natural than using employment, as is sometimes done. If shocks take time to propagate to the labor market, using output will correctly time business cycles at the state level. Additionally, using employment to measure business cycles presents problems when examining the effects on employment as an outcome, since there may be a mechanical relation even in the absence of an economic relation. Using output does not present this problem as it is measured by a distinct data source (see data description below).

Potential Endogeneity: The coefficient $\beta^{(k)}$ gives the response of y to a one-time, temporary, one percent shock to gross state product (GSP) k periods later. For $\beta^{(k)}$ to identify a causal effect of the GSP shock on $y_{s,t+k} - y_{s,t-1}$, it must be the case that, conditional on the set of controls, the growth rate of GSP in period t is uncorrelated with the error term:

$$\mathbb{E}\left[\Delta GSP_{s,t} \cdot \epsilon_{s,t+k} | X_{s,t}\right] = 0$$

However, two key concerns suggest that requirement might not be met in practice. One concern is that employment may affect GSP, as lower employment (either through higher unemployment rates or lower LFPRs or both) mechanically lowers GSP. A second concern is that GSP growth could be autocorrelated, in which case estimates of $\beta^{(k)}$ may pick up the correlation between $y_{s,t+k} - y_{s,t-1}$ and GSP growth rates in future (or past) periods.

Instrument: To overcome these issues, we instrument for ΔGSP with a Bartik (1991) shift-share type measure. The first-stage equation is as follows,

$$\Delta GSP_{s,t} = \alpha Bartik_{s,t} + \gamma X_{s,t} + \nu_{s,t} \tag{2}$$

where

$$Bartik_{s,t} \equiv \sum_q \Delta GDP_{q,-s,t} \omega_{q,s,t-5}. \tag{3}$$

Industries are indexed by q, and $\omega_{q,s,t-5}$ represents industry q's share of total GSP in state s five years previously. $\Delta GDP_{q,-s,t}$ represents the national gross domestic product growth in industry, q, for period, t using the "leave-one-out" approach–that is, we calculate $GDP_{q,-s,t}$ by summing up $GSP_{q,s,t}$ across all states except for state s.

Identifying Assumptions: In order for $Bartik_{s,t}$ to be a valid instrument, it must meet the following conditions (Stock and Watson, 2018):

$$
\begin{array}{lll}
(1) & \mathbb{E}\left[Bartik_{s,t} \cdot \Delta GSP_{s,t} | X_{s,t}\right] = \alpha \neq 0 & \text{(relevance)} \\
(2) & \mathbb{E}\left[Bartik_{s,t} \cdot \epsilon_{s,t} | X_{s,t}\right] = 0 & \text{(contemporaneous exogeneity)} \\
(3) & \left.\begin{array}{l} \mathbb{E}\left[Bartik_{s,t} \cdot \epsilon_{s,t+k} | X_{s,t}\right] = 0 \\ \mathbb{E}\left[Bartik_{s,t} \cdot \Delta GSP_{s,t+k} | X_{s,t}\right] = 0 \end{array}\right\} \text{for } k \neq 0 & \text{(lead-lag exogeneity)}
\end{array}
$$

$Bartik_{s,t}$ captures predicted GSP growth for a given state, s, in time, t, based on that state's industry mix in period $t-5$. We argue that this is likely to be relevant as local output in a given industry is likely to be correlated with national output in that industry due to changes in industry technology. The contemporaneous exogeneity assumption will hold as long as the national industry shocks used to construct $Bartik_{s,t}$ are unrelated to local changes in labor market outcomes (where we have removed any mechanical correlation by using a "leave-one-out" approach). Lead-lag exogeneity requires not only that $Bartik_{s,t}$ is uncorrelated with unobserved forces affecting local labor markets in other periods, but also that it is not correlated with either of the two components of $\Delta GSP_{s,t+k}$ in other periods.

2.2 Inference

This section describes three important issues for inference in our research design: the role of clustering in computing standard errors, testing for potential weak instruments, and how we weight observations.

Clustering: To quantify the uncertainty around our estimated impulse response functions, we compute heteroskedasticity-robust standard errors clustered at the state-level in our baseline specification. Adão, Kolesár, and Morales (2018) raise concerns that this approach may understate uncertainty in shift-share designs; however, our instrument is likely to be one for which state-clustered standard errors are appropriate.[1]

Testing for Weak Instruments: To verify that our estimates are not affected by weak instrument problems, we conduct first-stage F-tests for each specification. We compute the first-stage F-statistics under the assumption of homoskedasticity and examine whether they exceed 10 to determine if our instrument is weak, following Staiger and Stock (1997). Al-

[1] Our analysis is most similar to the results shown in Panel B of Table 6 in Adão, Kolesár, and Morales (2018), which shows that more sophisticated approaches to estimate confidence intervals are not meaningfully different from clustering by local labor market.

though the instrument and endogenous variable are the same in all specifications, the F-statistics may vary across regressions with different weights.

Weighting: We weight each regression of outcome $y_{s,t}$ for group j by the population, n^j_{st} of group j in state s at time t. Weighting has two main advantages in this setting. First, weighting the regressions by population allows us to interpret the estimates in terms of the national unemployment rate, LFPR, and EPOP. Second, the smallest states have relatively few respondents in the CPS, which has the potential to generate noise when calculating state-level unemployment rates, LFPRs, and EPOPs for those smaller states and yield imprecise regression estimates. Additionally, the noise issue compounds when slicing the data further into subgroups of the population, such as prime-age individuals, men and women, and levels of educational attainment. Weighting by state-level population reduces the role of noise in the estimates.

2.3 Data

We combine state-level data from multiple sources to form annual panel. In particular, labor market outcome variables consist of the unemployment rate, the labor force participation rate, and the employment-to-population ratio, which are all obtained from the Current Population Survey. Our output growth variable, $\Delta GSP_{s,t}$, is the BEA real gross state product (GSP), while the instrument $Bartik_{s,t}$ is based on BEA GSP data by industry. Finally, the control variables, $X_{s,t}$, include home price growth (FHA) as well as state and year fixed effects.

3 Results

This section presents estimates of the cyclicality of the unemployment rate, LFPR, and EPOP from estimating impulse response functions using our research design. After showing evidence that each of our outcomes do respond to cyclical shocks, we examine in detail how the responses of the outcomes compare to each other. We show that the LFPR responds with a substantial lag compared to the unemployment rate, although with a smaller amplitude, and remains persistently lower even well after the unemployment rate has fully recovered. This persistent effect is a consequence of changes in the demographic composition of the population induced by cyclical shocks, rather than a form of "hysteresis" affecting participation rates conditional on demographics. We also explore the heterogeneity of these responses across demographic subgroups, finding varying responses over the lifecycle and a sharp difference in responses between different education levels.

3.1 Cyclicality of Labor Market Outcomes

We start by examining the cyclical responses of labor market outcomes for the CPS population, which includes non-institutionalized adult civilians ages 16 and over. We use our local projections research design to estimate impulse response functions for the unemployment rate, LFPR, and EPOP from 3 years before the shock to 15 years after the shock. For ease of interpretation, we report all of our estimates as the response to a temporary 1 percent negative shock to GSP, so that the cyclical responses will have the same sign as in a recession.

Figure 1: Cyclicality for Ages 16 and Older

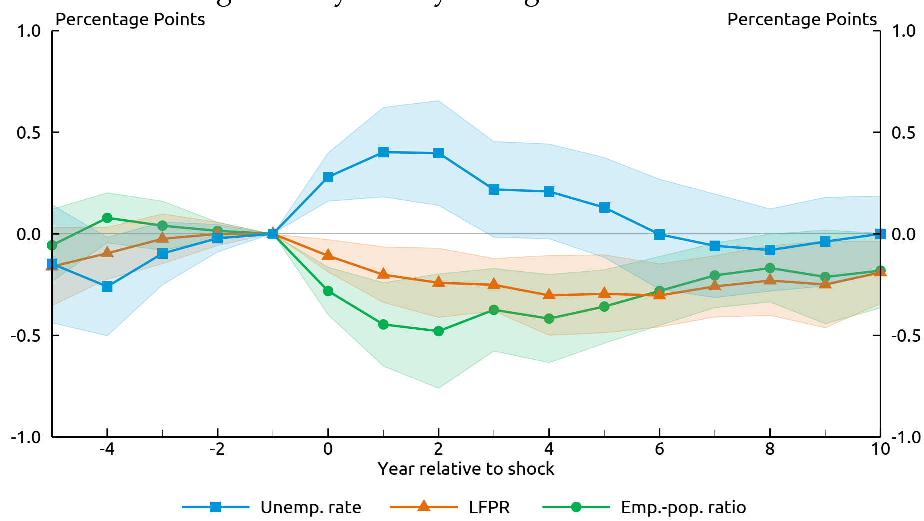

Note: F-statistic: 159.3. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, and authors' calculations.

The results show that the unemployment rate, LFPR, and EPOP all respond to cyclical shocks, but the timing of these responses varies considerably. Starting with the unemployment rate, a contractionary 1 percent shock to output causes contemporaneous increase in the unemployment rate of 0.3 percentage points. The increase in the unemployment rate continues in the following year and peaks at 0.45 percentage points one year after the shock. Our estimate of the total increase in the unemployment rate due to a negative one percent shock to GSP is squarely in the range of -0.5 to -0.4 of Okun's law coefficients estimated in the literature; see, for example, Balleer, Gomez-Salvador, and Turunen (2014). Following the peak one year after the shock, the unemployment rate steadily declines by about 0.1 percentage points per year until it returns to its pre-shock value about 6 years after the shock and remains there. The pattern of our estimated response of the unemployment rate to a GSP shock reflects the asymmetric rates of increase and decrease around the peak of the unemployment rate observed during and following a recession.

The LFPR also shows a significant response to a negative GSP shock, but with a substantial delay compared to the unemployment rate. For example, the LFPR declines by less than 0.1 percentage points contemporaneous with the shock, much less than the unemployment rate increases. However, as the unemployment rate increases for only one more year after the shock, the LFPR continues to steadily decline for several years after the shock, finally reaching a trough 4 to 6 years later at a level that is 0.3 percentage points below its initial value. After reaching its trough, the point estimate of the LFPR response recovers a bit, but it never fully recovers. Instead, it settles at a new long-run level about 8 years after the shock that is about 0.25 percentage points below its initial value.

The response of the EPOP to the GSP shock, of course, reflects the combination of the responses of the unemployment rate and LFPR to the shock. The EPOP declines by 0.3 percentage points at the onset of the shock, largely reflecting the initial spike in the unemployment rate. By two years after the shock, the EPOP reaches its trough at 0.5 percentage points below its pre-shock value, reflecting an unemployment rate that has peaked and an LFPR that is still declining. After reaching its trough, the EPOP begins a steady recovery, as the declining unemployment rate outweighs the still declining LFPR. However, at 8 years after the shock, the EPOP stops recovering and settles at a new long-run value about 0.2 percentage points below its pre-shock value. The persistently lower level of EPOP following the shock is entirely due to the LFPR, as the unemployment rate fully recovers but the LFPR does not.

3.2 The Role of Changing Demographic Composition

The persistent effect of GSP shocks on LFPR could represent either hysteresis or compositional shifts. By hysteresis we mean that individuals within a given demographic group become persistently less likely to participate in the labor market as a result of the shock. The alternative is that the composition of the population shifts in response to the shock such that groups with lower participation rates make up more of the population after the shock. We examine the role of composition changes directly, and treat any part of the persistent decline not explained by composition as coming from hysteresis effects.

To investigate whether the persistently lower levels of LFPR and EPOP in response to the shock are due to changes in the composition of the labor force, we decompose the labor market variables into the component explained by s and the component that is not. This analysis has multiple steps. First, using individual-level data from the CPS, we regress whether an individual was unemployed, employed, or participated in the labor market on that individual's demographic characteristics using the following linear-probability model:

$$Y_{i,s,m,t} = \psi_0 + \Psi^{i,m,t} D_{i,m,t} + \Psi^{s,m,t} X_{s,m,t} + \eta_{i,s,t} \tag{4}$$

where $Y_{i,s,m,t}$ is a dummy variable indicating whether individual i in state s was either employed, unemployed, or participating in the labor force in month m of year t; $D_{i,m,t}$ is a vector of dummy variables over the demographic characteristics of individual i in month m of year t that include age, gender, educational attainment, race/ethnicity, and marital status; and $X_{s,m,t}$ is a vector of state, month, year fixed effects.[2] We include month-of-year dummy variables to account for seasonality.

Figure 2: Cyclicality for Ages 16 and Older; Demographic and Demographically-Adjusted Components

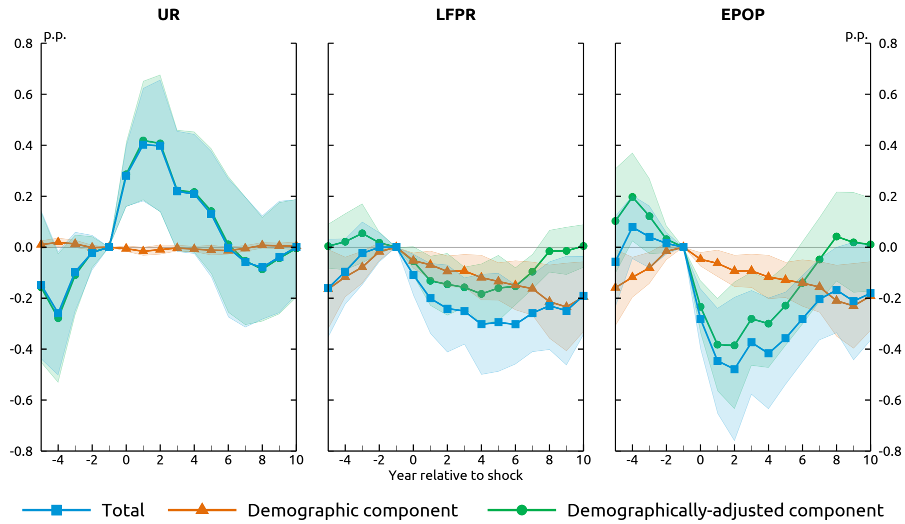

Note: F-statistic: 159.3. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

Using the estimated coefficients from equation 4, we predict whether an individual is unemployed, employed, or participating in the labor market based on the demographic characteristics in that equation, $\widehat{Y^{D}_{i,s,m,t}}$, and we use the residual of equation 4 as the demographically-adjusted component, $\widehat{Y^{D.adj}_{i,s,m,t}}$. We then aggregate the individual-level predicted and residual components to calculate monthly rates for the demographics-only and demographically-adjusted labor market variables in each state s, and then average across months within year t to create a demographics component, $\widehat{y^{D}_{s,t}}$, and a demographically-adjusted component, $\widehat{y^{D.adj}_{s,t}}$, for each of our labor market variables. Finally, we use those demographics-only and demographically-adjusted state-level variables as the dependent variable in equation 1.

[2]The age variables are single-year age dummies for ages 16 to 79 and a dummy variable for ages 80 years and older. The educational attainment dummies partition attainment into five categories: less than a high school degree, a high school degree, some college, a college degree, and more than a college degree. The race/ethnicity dummies partition the population into four groups: non-hispanic white, black, non-black hispanic, and other. Marital status is a single dummy indicating whether an individual is married.

Changes in the demographic composition of the labor force explain little of the unemployment response to cyclical shocks. As shown in Figure 2, the demographics component is flat before and after the shock to output. As a result, the entire response of unemployment is accounted for the demographically-adjusted component. Similarly, demographics account for little of the short-run response of LFPR and EPOP, as the majority of the response in the first few years for these variables is due to the demographically-adjusted component.

However, the demographic component accounts for nearly the entire long-run responses of LFPR and EPOP. The GSP shock causes the demographic components of LFPR and EPOP to steadily decline for many years after the shock, never showing signs of a sustained recovery and remaining permanently below their pre-shock value. In the first five years after the shock, the demographic components of both LFPR and EPOP decline steadily and significantly but decline less than the demographically-adjusted components. Beyond five years after the shock, though, the demograhic components continue to decline while the demographically-adjusted components recover. By eight years after the shock, the demographic components explain all of the decline in the LFPR and EPOP, as both remain below their pre-shock value by about 0.2 percentage points.

Accordingly, the demographically-adjusted components of the LFPR and EPOP, which represent

demographically-adjusted rates of these variables, fully recover after cyclical shocks. The GSP shocks do cause the demographically-adjusted components of LFPR and EPOP to decline, but the demographically-adjusted components of both variables eventually fully recover from the shock, although the full recovery occurs well after the initial shock and significantly later than the unemployment rate recovers. As shown in Figure 2, the demographically-adjusted component of the LFPR declines each year until it reaches its trough four years after the shock at 0.2 percentage points below its pre-shock level. After reaching its trough, the demographically-adjusted component of LFPR begins a steady recovery and, importantly, eventually fully recovers, though not until at least 8 years following the initial shock. The demographically-adjusted component of EPOP displays a similar pattern as the LFPR, though its decline on impact is much steeper than the LFPR and recovers slightly faster than the LFPR, reflecting the dynamics of the unemployment rate.

To better understand why the demographics-only components of LFPR and EPOP never fully recovery, we further decompose these components into the contributions of specifc demographic variables–that is, we calculate a demographic component attributed to the age composition only, a component attributed to the gender composition, and so on. As shown in Figure 3, the age-only component explains nearly all of the decline in the demogrpahics

Figure 3: Cyclicality for Ages 16 and Older; Breaking Down the Demographics Component

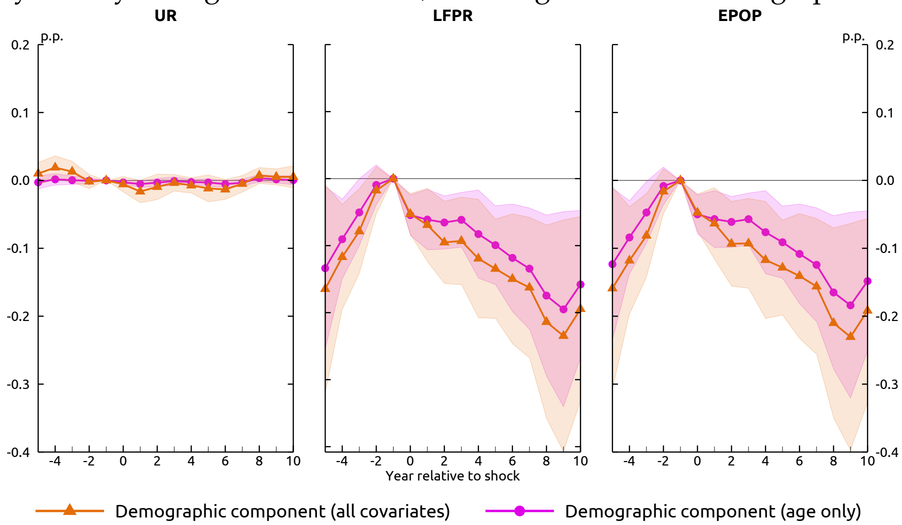

Note: F-statistic: 159.3. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

components of LFPR and EPOP in response to the shock.[3] Thus, the shock to GSP changes the age composition of the labor force in states exposed to the shock in way that persistently lowers the LFPR and EPOP. We will further explore why the age composition of the labor force persistently changes in response to the shock and the implications for measuring the cyclicality of the aggregate LFPR and EPOP below in section 4.

3.3 Examining Individuals Ages 25–54

The labor market performance of individuals between the ages of 25 and 54 (prime-age people) is often used as a benchmark for the cyclical state of the labor market as a whole, as changes in demographics (such as the the aging of the baby boomers into their retirement years) usually affect this group considerably less than the overall population. Although our results above in section 3.1 for the overall population control for these changes in demographics, and thus gives us a clean reading on the cyclical response of the labor market, understanding the cyclical response for the prime-age group is still of considerable interest, as the prime-age group makes us about 50 percent of the 16 and over civilian non-institutional population and roughly 60 percent of the labor force. Further, much work has focused on the structural factors contributing to the long-run and steady decline of the trend prime-age LFPR and EPOP (see, for example, (Abraham and Kearney, 2018)), but there has been relatively less work on identifying the cyclical response of those variables from their long-run

[3]The remaining part of the demographic component that is not accounted for by age alone is mostly attributable to education, with each of the other demographic variables contributing little on net.

declining trends.[4]

Figure 4: Cyclicality for Ages 25 to 54

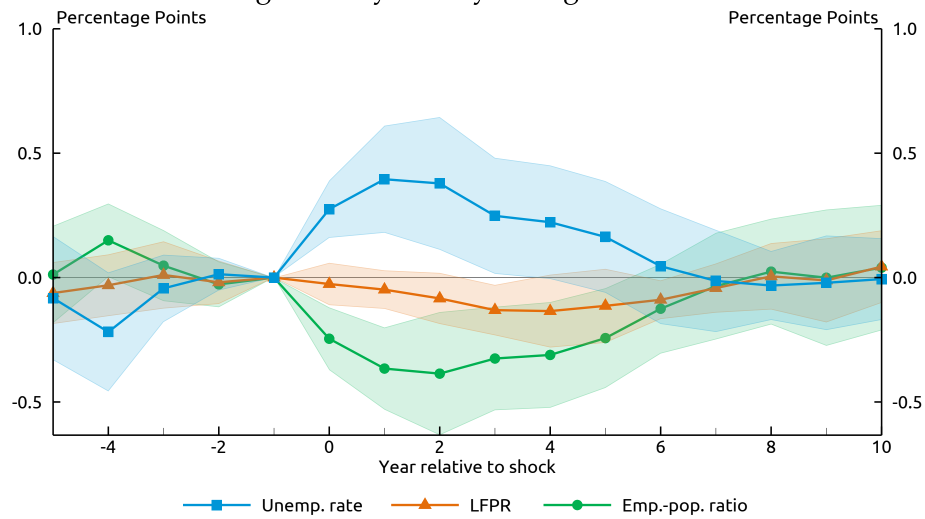

Note: F-statistic: 151.6. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

The cyclical response of the prime-age unemployment rate is similar in both timing and magnitude compared to the overall unemployment rate. Indeed, as shown in figure 4, the unemployment rate peaks one year after the GSP shock at around 0.4 percentage points, then steadily declines until it fully recovers at 7 years after the shock. The prime-age unemployment rate completes its recovery one year later than the overall unemployment rate, but the standard error for the prime-age estimate is a touch larger compared to the overall population. We cannot rule out that the timing of recoveries for the overall and prime-age unemployment rates are not the same.

The cyclical response of the prime-age LFPR and EPOP are also similar to the overall population's response for those variables in the initial response and decline, but the recovery patterns for both variables are considerably different across the two groups. The LFPR declines steadily after the shock until it reaches its trough 3 to 4 years after the initial shock—well after the unemployment rate peaks—at about .2 percentage points below its pre-shock level, whereas the EPOP reaches its trough two years after the shock at about .45 percentage points below its pre-shock level. Like the unemployment rate, the prime-age LFPR and EPOP estimates have wider confidence intervals than the estimates for the overall population. In the case of the prime-age LFPR, the LFPR is statistically lower than its pre-shock

[4]Although the main purpose of (Aaronson, Cajner, Fallick, Galbis-Reig, Smith, and Wascher, 2014) and (Montes, 2018) is to build a forecasting model of the overall, 16 years and older LFPR, both papers provide some evidence on the cyclicality of prime-age LFPR. Our work complements those papers in that we establish a causal response to output shocks, whereas those estimates were largely based on correlations with changes in the unemployment rate.

level only in the third year after the shock.

Whereas the overall LFPR and EPOP remained persistently below their pre-shock values in the long-run, the prime-age LFPR and EPOP both steadily increase after reaching their troughs and fully recover about 8 years after the shock. Thus, the shock does not cause a change in the age composition within the prime-age population in a way that persistently lowers the prime-age LFPR, even though it does change the age composition of the overall population in a way that lowers the overall LFPR. These patterns suggest that the change in the age composition of the population in response to the shock likely reflects a shift in population shares away from prime-age people and towards younger and older people. We will formally address this composition shift in the population due to the shock in section 4.

Figure 5: Cyclicality for Ages 25 to 54 by Sex

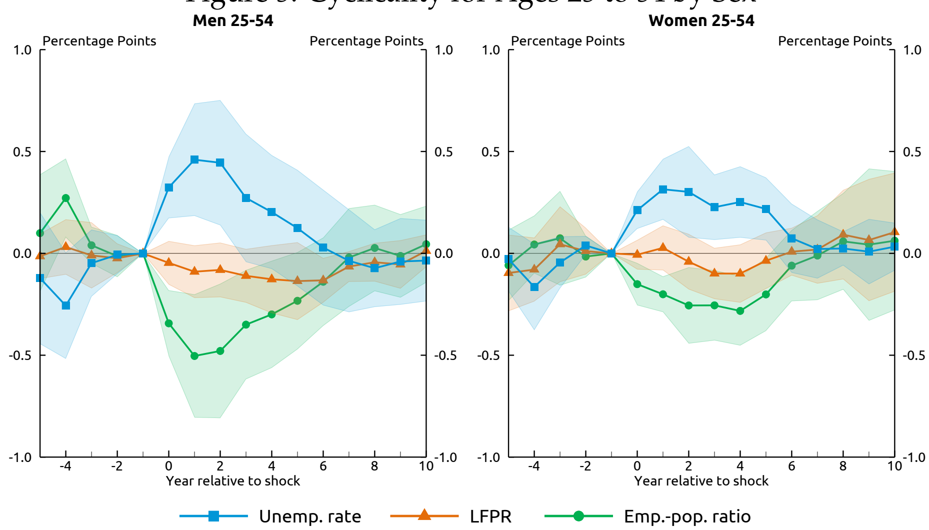

Note: F-statistic: 151.1 for men, 152.1 for women. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

Digging deeper into the prime-age labor market responses, our results suggest that while both men and women have strong cyclical responses to a GSP shock, the magnitudes and timing of their responses are quite different. The unemployment rate spikes for men at the impact of the shock and climbs further to an increase of 0.5 percentage points a year after the shock, as shown in the left panel of Figure 5. The unemployment rate for prime-age men then steadily declines in each subsequent year until fully recovering 7 years after the shock. Prime-age women experience an increase in their unemployment rates on impact and a year after the shock, just as men, but the increase is smaller, reaching about 0.3 percentage points (right panel of Figure 5). The recovery pattern for prime-age women is different than for men, though, as their unemployment does not experience a steady and continuous decline. Rather, it remains elevated at about 0.25 percentage points above its pre-shock value as many as five years after the shock and only recovers about 7 to 8 years after the initial shock. That the unemployment rate for prime-age women remains near its peak for so many years after the shock, though, is a pattern observed in the data. For example, after both the 2001 and 2007-2009 recession, the women's prime-age unemployment rate did not start to recover until at least 3 years after the end of the recession.

The response of prime-age LFPRs to the shock are also quite different across gender. For men, the initial point estimate response is small, and subsequent year-over-year declines are also small. However, even though those yearly declines are small, they compound for many years after the shock, cumulating to a total decline in the LFPR of about .15 percentage points at its trough 6 years after the shock. Although the confidence bands around those estimates are large due the smaller sample sizes from splitting the prime-age group by gender, the decline the prime-age LFPR for men is large enough in year 6 for the confidence band to not include zero.

The response of LFPR for prime-age women is considerably delayed. In fact, the LFPR of prime-age women does not start to decline until 2 years after the shock and reaches its trough 3 to 4 years after the shock at about 0.1 percentage points below its initial value. This rate fully recovers by about 6 years after the shock and settles at rate slightly above its pre-shock value. Of course, the confidence bands around the estimates for prime-age women are quite large, possibly due to large non-cyclical variation in the LFPR for prime-age women, and so one cannot reject the possibility that the LFPR of prime-age women does not respond to the shock at all.

All told, our point estimates suggest that the prime-age labor market does not fully until 8 years after the initial GSP shock. Of course, there is uncertainty around that estimate, but the lower end of the 95-percent confidence band still suggest that the earliest a fully recovery is prime-age EPOP occurs is 6 years after the shock, and the upper end of the confidence band suggests that the recovery could still be taking place more than 8 years after the shock.

3.4 Other Age Groups

Although the prime-age population often receives the most focus when discussing the health of the labor market, younger people (ages 16 to 24) and older people (ages 55 and over) comprise roughly half of the overall population and 35 percent of the labor force. Moreover, those shares have slowly but steadily been increasing over the past 20 years, and are likely to continue increase into the near future, as the baby boomers continue to age in their retirement years. As a result, the responses of younger and older workers to a recession will have a large effect on how the labor market variables for the overall population respond an output shock.

The cyclicality of younger and older workers' labor market outcomes is especially difficult to measure from aggregate data. Decomposing changes in the LFPRs and EPOPs of younger and older people into changes caused by a recession as opposed to changes reflecting long-run trends is difficult, as the LFPR for both groups has trended sharply over the past 30 years.[5] For younger workers, the 12 percentage point decline in their LFPR reflects both an increase in school enrollment rates and a decrease in the LFPR of those enrolled in school, and these trends are likely a response to the high and increasing college wage premium over that period. For older workers, the 10 percentage point increase in their LFPR likely reflects a combination of changes in the age composition within the 55 and over group, an increase in the health capacity to work at older ages, and increases in the age at which individual's can collect full retirement benefits through Social Security.[6] Since these factors are likely to affect all states, our approach of leveraging business cycle shocks across states controls for these trends and allows us to identify the cyclical response of these groups.

Figure 6: Cyclicality for Ages 16-24 and 55+

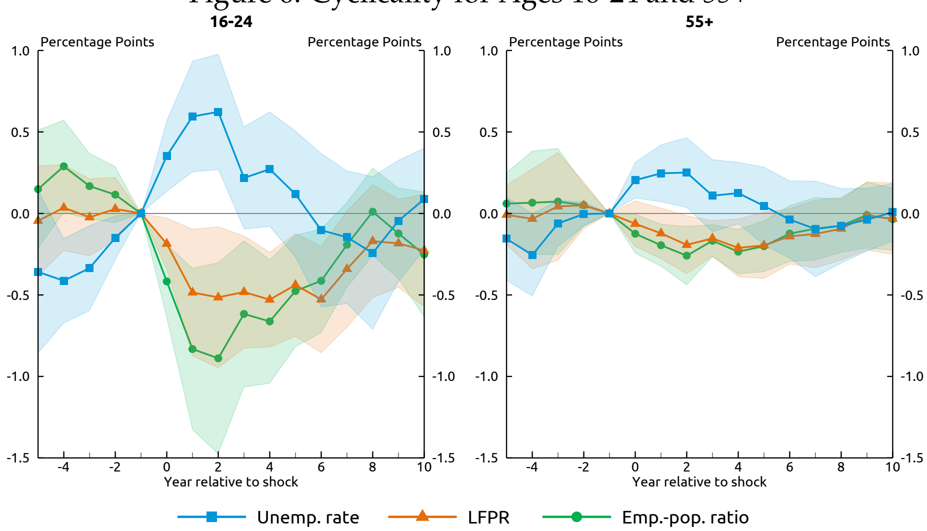

Note: F-statistic: 179.1 for 16–24, 162.6 for 55+. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

The unemployment rate for younger people has a larger cyclical response than the overall population, but recovers at a slightly faster pace. The unemployment rate for younger people is about 0.6 percentage points higher than the pre-shock level by 1 year after the shock, a slightly larger increase compared to the unemployment response of prime-age peo-

[5]The sharp trends in EPOP for both younger and older people over the past 30 years entirely reflects their sharply trending LFPRs, as the unemployment rates for both groups show no clear, long-run trends.

[6]For general discussions on these factors, see (Aaronson, Cajner, Fallick, Galbis-Reig, Smith, and Wascher, 2014), (Montes, 2018), and (Bauer, Moss, Nunn, and Shambaugh, 2019).

ple and the overall population. However, this rate recovers 5 to 6 years after a shock, slightly faster than the 6 to 7 year recovery of the prime-age unemployment rate, although the confidence bands are large enough around both estimates as to not rule out that both unemployment rates recover concurrently.

The recovery of LFPR for younger people is more delayed than the overall population, even though it reaches its trough sooner than the LFPR for the overall population. The LFPR for young people reaches its trough, more or less, one year after the shock—three years before the overall LFPR reaches its trough—at a level of 0.5 percentage points below its pre-shock value. Rather than beginning its recovery soon after reaching its trough, the LFPR for younger people lingers near at its trough level for an additional five years and does not begin its recovery until 7 years after the shock—2 years after the overall LFPR starts its recovery. The point estimate of the LFPR of younger people never fully recovers, as it settles at about 0.2 percentage points below its pre-shock value, although the upper end of the confidence interval suggests we cannot rule out a full recovery.

The unemployment rate and LFPR response for older people to the shock is qualitatively similar to the response of younger people, but the magnitude of the response of older people is much more muted. The unemployment rate for older people peaks at 2 years before fully recovering 5 to 6 years after the shock, just as it does for younger people; however, its peak value is about 0.25 percentage points above its pre-shock level—much lower than the peak for younger people. Similarly, the magnitude of the decline in the LFPR of older people is a little less than half that of younger people, reaching its trough at about 0.2 percentage points. The LFPR for older people then lingers around its trough for about 3 to 4 years before beginning a sustained recovery, again similar to the LFPR for younger people. Unlike with younger people, though, the LFPR for older people eventually recovers 9 years after the shock, although the point estimate then jogs down a bit after that recovery.

3.5 Differences in Cyclicality by Education

Labor market outcomes over at least the past 40 years have been quite different for lower- and higher-educated individuals. Indeed, the levels of the unemployment rates, LFPRs, and EPOPs for prime-age workers vary significantly across levels of educational attainment for both men and women. Additionally, the prime-age LFPR and EPOP for lower-educated people have been declining steadily over the past several decades, while the LFPR and EPOP for higher-educated prime-age people were relatively flat. Those trends have led to a growing divergence in labor market outcomes between the most and least educated individuals.

This divergence may, at least in part, be due to a long-term decline in the demand for lower-educated workers that is unrelated to the business cycle and caused, perhaps, by changes in technology and globalization; thus, we need to control for the long-term structural declines to identify the response to a recession. As was the case for older and younger people, we isolate the effects of the business cycle and control for long-term trends by taking advantage of the variation from state-level business cycles. Examining state-level recessions allows us to control for national- and state-level trends and leverages the different severity of business cycles across states to identify the typical patterns of labor market behavior across educational groups.

Figure 7: Cyclicality for Ages 25 to 54 by Education

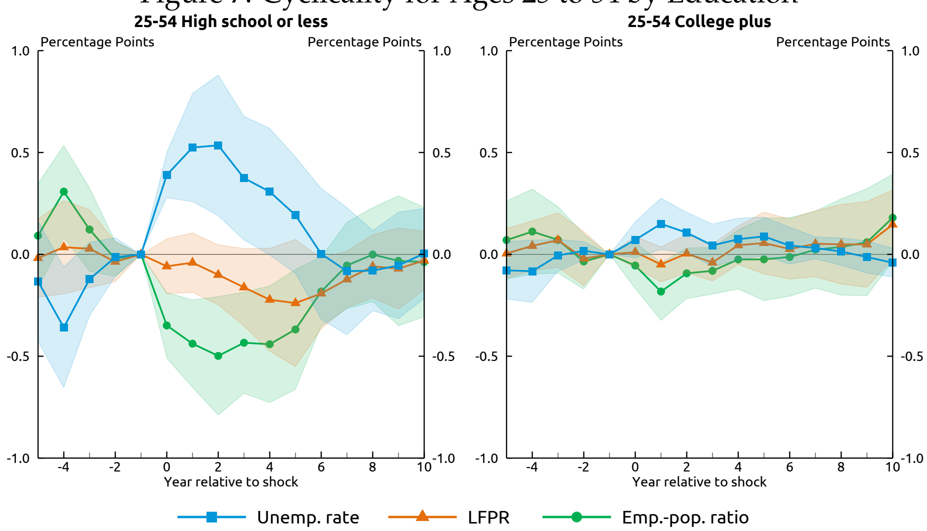

Note: F-statistic: 167.9 for high school degree or less, 137.7 for college degree or more. Regressions are weighted by population. Standard errors clustered by state.
Source: BLS, BEA, own calculations.

We find a starkly different evolution of labor market over the business cycle for less-educated prime-age workers compared to those with college degrees even after controlling for national and state trends. In response to the 1 percent decline in output growth that returns to normal in the next year, the unemployment rate increases and the EPOP decreases immediately for both groups. However, those responses are considerably larger for prime-age workers with a high school degree or less, as the total increase in the unemployment and decline in the EPOP is about 4 times as large for the less-educated prime-age workers than it is for the more-educated prime-age workers (see figure 7).[7]

The unemployment rate and EPOP responses are also longer lasting for less-educated individuals. Part of this difference is due to the larger initial response of the less-educated unemployment rate and EPOP, but part is also due to a delayed start to the recovery for

[7]We omit workers with some college but less than a four year degree for ease of comparison. The labor market response of this group falls in between the two groups shown here, closer to the less-educated group than to the more-educated group.

the less-educated group. For example, the unemployment rate for less-educated workers does not start its recovery until three years after the shock, whereas their EPOP does not begin to recover in earnest until about 5 years after the shock—significantly later than the more-educated group. Once the less-educated group's unemployment rate and EPOP begin a sustained recovery, though, both rates eventually fully recover, but not until well after the initial shock. In particular, the EPOP of the less-educated group does not fully recover until 8 years after the initial shock, whereas the point estimate of the more-educated group's EPOP is within a few basis points of its pre-shock value after 4 years and fully recovered after 6 years.

A large portion of the delayed recovery in the EPOP of the less-educated workers is due to a long delay in the response to and recovery from the shock of this group's LFPR. Upon impact of the shock, the LFPR for less-educated workers barely falls at all. In fact, this group's LFPR does start to steadily decline until 2 years after the shock, vastly different than the response of this group's unemployment rate, which experiences essentially all of its increase within the first year of the shock. Additionally, once the LFPR of the less-educated starts to decline, it does not stop until it reaches its trough 5 years after its shock. From that point, this group's LFPR recovery is slow, as its point-estimate only approaches its pre-shock level 8 to 10 years after the shock and does not reach its pre-shock level until 12 years after the contraction.

4 Response of Population Composition to Cyclical Shocks

Our results show that the LFPR is highly cyclical, but responds to and recovers from a business cycle shock with a significant delay, a robust pattern across many subgroups. Indeed, for each subgroup, the LFPR declines slowly but steadily each year for 4 to 5 years following a shock, reaches its trough, and then slowly but steadily increases before fully recovering between 8 and 10 years after the initial shock.

The one group that breaks from that pattern, though, is the LFPR response for the 16 and over population as a whole. The LFPR for the overall population declines for about 4 years after the shock and then remains near its trough in perpetuity, never recovering. As we show in section 3.1, the persistently lower level of LFPR in response to the business cycle shock is entirely explained by demographics, and, in particular, a change in the age-composition of the population in response to the shock. The demographically-adjusted component of the overall population's LFPR, though, responds to a business cycle shock similarly to how the subgroups respond: steadily declining for about 4 to 5 years following the shock and fully recovering about 8 to 10 years after the shock. Taken together, these patterns suggest state-level LFPRs for the overall population are persistently lower because, in response to the shock, the age-composition of the population in a given state shifts away from groups that participate at higher rates towards groups that participate at lower rates; the LFPRs of these various groups, though, respond similarly to a business cycle shock and eventually fully recover.

Why does the age-composition of the population change in response to a business cycle shock? Blanchard and Katz (1992) provide empirical evidence that economic shocks at the state level trigger adjustments not only through unemployment, but also by triggering cross-state migration. More recently, (Dao, Furceri, and Loungani, 2017) show that it still remains the case that net migration across states responds to spatial disparities in labor market conditions and especially so during recessions, though the effect has weakened somewhat over time. However, both papers estimate the response of total net migration across states, and neither paper shows whether the response of net migration is concentrated among specific subgroups. The response of migration among specific subgroups may matter, even holding the total response constant; for example, if a business cycle shock triggers a permanent net out-migration of prime-age individuals (who tend to have higher LFPRs) relative to non-prime-age individuals (who tend to have lower LFPRs), then the overall LFPR of a state hit by a business cycle shock will be permanently lower.

Investigating the migration channel is crucial for interpreting our state-level results in the context of the aggregate economy. Inferring the cyclicality of the aggregate LFPR and EPOP from our state-level regressions can been problematic if we do not account for migration. For example, if prime-age individuals (who have above-average LFPR) move from a state hit by a negative shock to a state not hit by a shock, then this will lower (raise) LFPR in the first (second) state, but those migration patterns may leave the LFPR in the U.S. as a whole unaffected. In that case, controlling for changes in the age composition of the population across states—as we do in section 3 by estimating the response of LFPR and EPOP coming from the demographics-only and demographically-adjusted components—is necessary to recover the true cyclical response. In this section, we examine how the age composition of a state's population responds to a business cycle shock.

4.1 Results by Single-Year Age Group

To estimate the effect of a business-cycle shock on the composition of the state's population, we estimate the local projections equation (1) with $y_{s,t+k}$ being the log population of a single-year-age group in state s in period $t + k$. We estimate this equation for each single-year-age group from ages 16 through 80. The interpretation of the estimated equation for single-year-age group 25 in period $k = 10$ would be, for example, the percent change in the level of the total 25 year old population in state s between periods $t + 10$ and $t - 1$ caused by the business-cycle shock.[8] Thus, our work in this section adds to the work by Blanchard and Katz (1992) and Dao, Furceri, and Loungani (2017) by estimating the population effects of single-year-age groups and identifying whether compositional effects—particularly, younger prime-age individuals—are driving the cyclical net migration results that those papers find.

Figure 8: Percent Change in Single-Age Population in Response to a Business Cycle Shock

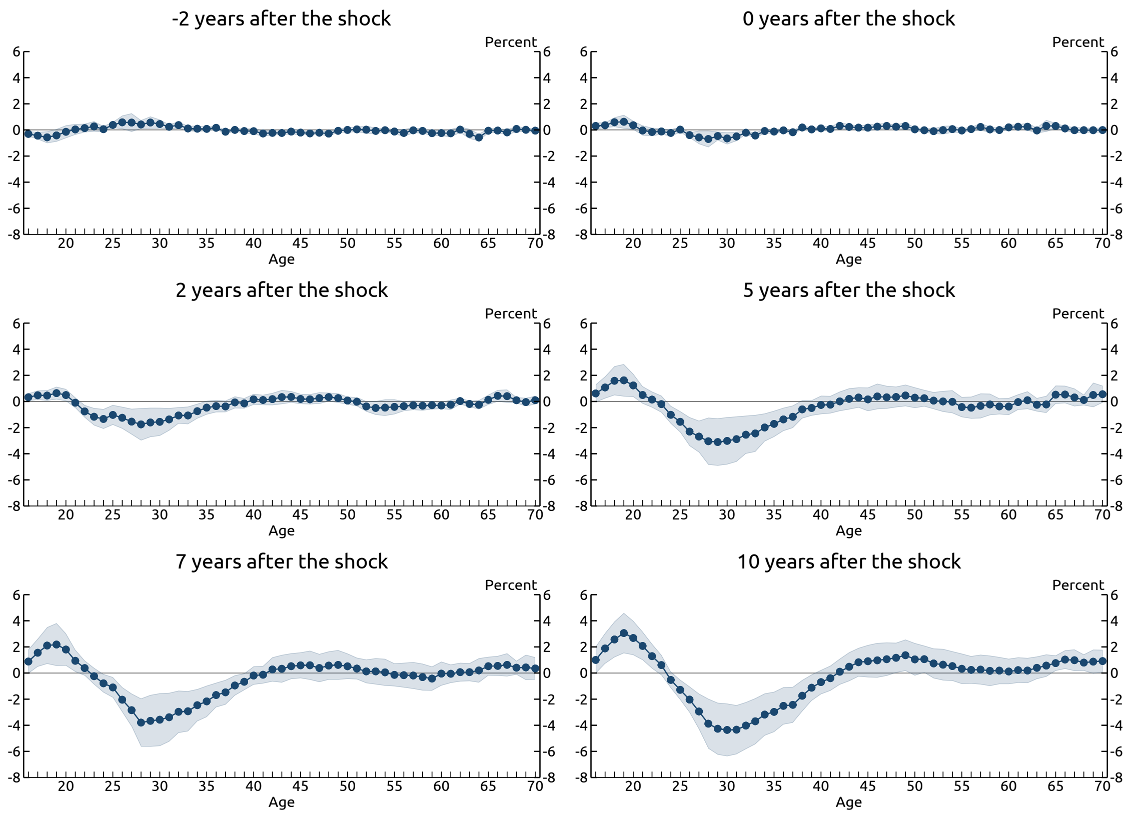

Note: The dependent variable is the percent change in the population of a single-age group in period $t + k$ relative to period $t - 1$. Regressions are weighted by population, and standard errors clustered by state.
Source: BLS, BEA, and authors' calculations.

A negative business cycle shock causes the population between the ages of 25 and 40 to

[8] We use state-level data for the covered-area population for single-year-age groups from the U.S. Census Bureau (https://www.census.gov/data/tables/2017/demo/popest/state-detail.html). These population estimates use the most recent decennial census population counts as a base and then add births, subtract deaths, and add net migration (both international and domestic) to produce yearly population estimates for each age in each state. The covered-area population is slightly different from the civilian noninstitutional population, which is used to calculate LFPR and EPOP. The main difference is that the covered-area population includes active members or the armed forces and those in institutions (e.g. penal, mental facilities, and homes for the aged), whereas the civilian noninstitutional population does not include these groups. This distinction is not likely to matter in our analysis.

persistently decline in states exposed to the shock relative to those states without a shock (see Figure 8, which shows the population response to a shock at -2, 0, 2, 5, 7, and 10 years after the shock.) Two years prior to the shock, there is limited evidence that changes in the population are correlated with the business cycle shock, as essentially all age groups have point estimates that are precisely estimated at 0.0 percent. Upon impact of the shock, the migration response is still small, with essentially all point estimates a precise 0.0. As time goes by, though, changes in the composition of the population become apparent. Two years after the shock, the population levels of 23 to 35 year olds are all about 2 percent below their levels immediately prior to the shock. Five years after the shock, the population of 27 to 33 year olds falls to 3 percent below its pre-shock level, whereas the populations of 24 to 26 year olds and 34 to 27 year olds are 2 percent below their pre-shock levels. Seven years after the shock, the population levels of 28 to 31 year olds fall to 4 percent below their pre-shock levels, whereas the population levels 24 to 27 year olds and 32 to 39 year olds are all significantly lower, ranging between 1 and 3 percent below their pre-shock levels. Ten years after the shock, the population levels of 29 to 31 year olds decline to about 5 percent below their pre-shock levels, and the population levels of all single-year-age groups between 25 and 39 years olds are at least 2 percent below their pre-shock values. Though not reported, the population responses 10 years after the shock tend to hold in years 11 through 15, suggesting that a negative business cycle shock permanently lowers the population of 25 to 39 year olds in exposed states.

This pattern suggests that the net out migration of a state's population caused by a negative business cycle shock is entirely driven by individuals between the ages of 25 and 39 years old. Since 25 to 39 year olds are among the highest in LFPRs relative to other age groups, permanent declines in a state's population that are concentrated in this age range will also permanently lower its LFPR through compositional effects, all else equal.

There are several plausible reasons why the out-migration response might be concentrated in individuals ages 25–39, although formally testing these theories is outside the scope of the present paper. First, individuals in this age range may be less likely to be homeowners, on average, so it might be easier for them to move to a different state in response to a negative shock. Additionally, if a state has been hit by a negative business cycle shock, individuals from others states that are finishing school may be less likely to move to such a state. As a result, if a state experiences a recession, it could have a "missing generation" of recent graduates. This is consistent with the responses shown in Figure 8, as initially, the largest response is for individuals in mid-20s. However, as time goes by (and individuals get older), the response shifts to the right of the age distribution.

5 Conclusion

We estimate the effect of a business cycle shock on the labor force participation rate and show that the LFPR is cyclical, but it responds with a smaller elasticity, a more delayed impact, and a more persistent effect than the unemployment rate. Our approach uses state-level variation in business cycles to estimate the cyclicality of LFPR and instruments for changes in state output with a shift-share instrument to establish a causal link between business cycle shocks and the dynamic response of LFPR. We estimate this dynamic response of LFPR to an output shock using the local projections regressions. This method is particularly well-suited for estimating LFPR's cyclicality and its lag structure compared to more traditional time series models, as its flexibility allows for the possibility of long-run effects of a business shock on LFPR, such as hysteresis, and does not impose strict assumption about the smoothness of trends—a particular concern for LFPR given the aging of the population and other longer-term structural change such as the inflow of women into the labor force.

We show that the LFPR declines only slightly upon impact of a business cycle shock but that it continues to steadily decline for at least four years following the shock, cumulating a considerable decline and troughing well below its pre-shock level. The dynamic behavior of the LFPR is in stark contrast to the response of the unemployment rate to the business cycle shock, which spikes upon impact, peaks shortly after the shock, and is nearly fully recovered to its pre-shock value by the time LFPR reaches its trough. Further, the LFPR shows only little signs on recovery and remains below its pre-shock level even 10 years after the initial shock, whereas the unemployment rate fully recovers by 6 years after the shock.

We find no evidence, though, that the persistently lower LFPR in states exposed to a shock reflects hysteresis; instead, the long-run decline in state-level LFPR is driven by a change in the composition of the population. Decomposing LFPR into a component explained by demographic characteristics—such as age, gender, educational attainment, race/ethnicity, and marital status—and a component that is adjusted for those demographic characteristics, all of the persistent decline in the LFPR is contained within the component explained by demographics, and in particular, changes in the age composition of the population. Indeed, we show that negative shocks to output lead to a net out migration for adults between the ages of 25 and 39 years old and that a permanently lower level of the population for these groups persists for more than 10 years after a business cycle shock. Thus, the persistently lower level of LFPR in states exposed to a shock is due to a shift in the composition of the population away from individuals in their prime working years, who have higher LFPRs, and towards younger and older workers, who have relatively lower LFPRs.

Since the shift the affected state's composition likely reflects individuals moving to better performing state, the persistently lower demographics component of the affected state-level LFPR are likely to cancel out in the aggregate, suggesting that the demographically-adjusted component provides a more pure read of the cyclical response of the LFPR in the aggregate.

The demographically-adjusted LFPR that controls for changes in the age composition (and the composition of other demographic variables) shows a clear cyclical pattern, with a clear decline and recovery pattern, though that pattern is considerably delayed relative to the shock. The response of the demographically-adjusted LFPR shows a slow but steady decline for the four years after the shock, just as the response of the actual LFPR. In contrast to the response of the actual LFPR, the demographically-adjusted LFPR begins a slow and steady recovery in the fifth year after the shock, while the actual LFPR stagnates. The demographically-adjusted LFPR eventually fully recovers to its pre-shock value—about 8 to 10 years after the initial shock. Moreover, a full recovery of the demographically-adjusted LFPR is supported by our results estimating the LFPR responses by various age and educational-attainment subgroups of the population. While those subgroups of the population display considerable heterogeneity in the depth of their LFPR's response to the shock, the number of years their LFPR declines after a shock, and the number of years their LFPR takes to fully recover to its pre-shock level, all of those groups display the pattern of a slow and delayed decline to the shock, a trough in their LFPR many years after the initial shock, and an eventual recovery to its pre-shock level that occurs years after their unemployment rate has recovered. These subgroup patterns are consistent with our results for the overall population where changes in composition drive the persistently lower LFPR in states exposed to a shock.

Taking all of these facts together, the LFPR is cyclical and does fully recover from a negative business cycle shock, but the LFPR's decline and eventual recovery is slow and occurs well after the initial shock. Further, this LFPR pattern has a characteristic of business cycles at least since 1980. Thus, studying the aggregate unemployment rate on its own without taking into account labor force participation and its diverse changes across groups offers an incomplete picture of the labor market response to business cycle shocks.

References

AARONSON, S., T. CAJNER, B. FALLICK, F. GALBIS-REIG, C. SMITH, AND W. WASCHER (2014): "Labor Force Participation: Recent Developments and Future Prospects," *Brookings Papers on Economic Activity*, 45(2 (Fall)), 197–275.

ABRAHAM, K. G., AND M. S. KEARNEY (2018): "Explaining the Decline in the U.S. Employment-to-Population Ratio: a Review of the Evidence," Working Paper 24333, National Bureau of Economic Research.

ADÃO, R., M. KOLESÁR, AND E. MORALES (2018): "Shift-Share Designs: Theory and Inference," Working Paper 24944, National Bureau of Economic Research.

AMIOR, M., AND A. MANNING (2018): "The Persistence of Local Joblessness," *American Economic Review*, 108(7), 1942–70.

BALAKRISHNAN, R., M. DAO, J. SOLE, AND J. ZOOK (2015): "Recent U.S. Labor Force Dynamics; Reversible or not?," IMF Working Papers 15/76, International Monetary Fund.

BALL, L., D. LEIGH, AND P. LOUNGANI (2017): "Okun's Law: Fit at 50?," *Journal of Money, Credit and Banking*, 49(7), 1413–1441.

BALLEER, A., R. GOMEZ-SALVADOR, AND J. TURUNEN (2014): "Labour Force Participation across Europe: a Cohort-Based Analysis," *Empirical Economics*, 46(4), 1385–1415.

BARTIK, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?*, no. wbsle in Books from Upjohn Press. W.E. Upjohn Institute for Employment Research.

BAUER, L., E. MOSS, R. NUNN, AND J. SHAMBAUGH (2019): "Employment, Education, and the Time Use of American Youth," Working paper, The Hamilton Project, Brookings.

BLANCHARD, O. J., AND L. F. KATZ (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 23(1), 1–76.

BORUSYAK, K., P. HULL, AND X. JARAVEL (2018): "Quasi-Experimental Shift-Share Research Designs," Working Paper 24997, National Bureau of Economic Research.

COUNCIL OF ECONOMIC ADVISERS (2014): "The Labor Force Participation Rate since 2007: Causes and Policy Implications," mimeo.

DAO, M., D. FURCERI, AND P. LOUNGANI (2017): "Regional Labor Market Adjustment in the United States: Trend and Cycle," *The Review of Economics and Statistics*, 99(2), 243–257.

ERCEG, C. J., AND A. T. LEVIN (2014): "Labor Force Participation and Monetary Policy in the Wake of the Great Recession," *Journal of Money, Credit and Banking*, 46(S2), 3–49.

GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2018): "Bartik Instruments: What, When, Why, and How," Working Paper 24408, National Bureau of Economic Research.

JORDÀ, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

MONTES, J. (2018): "CBO's Projection of Labor Force Participation Rates," Working Papers 2018-04, Congressional Budget Office.

STAIGER, D., AND J. H. STOCK (1997): "Instrument Variables Regression with Weak Instruments," *Econometrica*, 65(3), 557–586.

STOCK, J. H., AND M. W. WATSON (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, 128(610), 917–948.

YAGAN, D. (2019): "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 127(5), 2505–2558.