

The Long-Lived Cyclical-ity of the Labor Force Participation Rate*

Tomaz Cajner

John Coglianesi

Joshua Montes

July 6, 2021

Abstract

How cyclical is the U.S. labor force participation rate (LFPR)? We examine its response to exogenous state-level business cycle shocks, finding that the LFPR is highly cyclical, but with a significantly longer-lived response than the unemployment rate. The LFPR declines after a negative shock for about four years—well beyond when the unemployment rate has begun to recover—and takes about eight years to fully recover after the shock. The decline and recovery of the LFPR is largely driven by individuals with home and family responsibilities, as well as by younger individuals spending time in school. Our main specifications measure cyclical-ity from the response of the age-adjusted LFPR, and we show that it is problematic to use the unadjusted LFPR when estimating cyclical-ity because local shocks spur changes in the population of high-LFPR age groups through migration. LFPR cyclical-ity varies across groups, with larger and longer-lived responses among men, younger workers, less-educated workers, and Black workers.

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. We thank our discussants Bruce Fallick, Laura Giuliano, and Eliana Viviano, as well as Andrew Figura, Brendan Price, and seminar participants at the Federal Reserve Board of Governors, WALES 2020, SOLE 2020, UIUC, FAU Nuremberg, ASSA 2021, the 4th IZA Labor Statistics Workshop, and the Federal Reserve System Applied Microeconomics Conference for helpful comments and suggestions. 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.

For policymakers, the key question is: What portion of the decline in labor force participation reflects structural shifts and what portion reflects cyclical weakness in the labor market?

Janet Yellen (2014)

1 Introduction

How cyclical is the U.S. labor force participation rate? Many observers have noted that the labor force participation rate (LFPR)—the share of population 16 years or older that is either working or looking for work—exhibits some degree of cyclicity (see, for example, [Aaronson et al., 2014b](#); [Council of Economic Advisers, 2014](#); [Erceg and Levin, 2014](#); [Montes, 2018](#)). Measuring the degree of cyclicity in the LFPR is complicated, though, by the presence of trend movements reflecting structural changes in the labor market that are unrelated to the business cycle, including the prolific entry of women into the workforce through at least the 1990s and the aging of the baby boom generation since the late 1990s. Observers often disagree about the magnitudes of these trends, which results in substantial disagreement about the extent of cyclicity in labor force participation. Those disagreements can be particularly acute following recessions, such as the period following the Great Recession in which estimates of the cyclical portion of the LFPR shortfall varied from 20-60 percent ([Council of Economic Advisers, 2014](#)).

We estimate LFPR cyclicity using state-level business cycles, which sidesteps the need to identify trend changes in labor force participation at the national level. We use the local projections method introduced by [Jordà \(2005\)](#) to estimate the response of the state-level, age-adjusted LFPR to changes in state-level output. By using this approach, we are able to identify the response of the LFPR to unexpected declines in output without imposing strict parametric assumptions or assuming that the effects of business cycle shocks dissipate in the long run. To avoid endogeneity between output and the labor market, we instrument for changes in state-level output with a shift-share instrument exploiting variation in local exposure to national changes in output across industries ([Bartik, 1991](#)).

We show that labor force participation *is* cyclical, but that its response to an exogenous output shock is long-lived. In response to a negative 1 percentage point output growth shock, the LFPR declines slowly yet persistently and does not reach its trough until 4 years later—at about 0.2 percentage point below its initial value.

The LFPR then gradually recovers and eventually returns to its pre-shock level, but not until about 8 years after the initial shock.

The cyclical response of the LFPR substantially lags behind the unemployment rate. Following a negative 1 percentage point output growth shock, the unemployment rate spikes quickly and peaks a year later, with a peak response that is about 0.4 percentage point.¹ By the time that the unemployment rate fully recovers 6 years after the shock, the LFPR is still in the early stages of its cyclical recovery. The delay in recovery between the LFPR and the unemployment rate suggests that observers who focus only on the unemployment rate 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.

These results shed light on the extent of slack in the post-Great-Recession labor market, which was hotly debated by policymakers at the time. By 2014, the unemployment rate had nearly returned to its pre-recession level, but the LFPR had continued to decline, reaching about 3 percentage point below its pre-recession level. This led to substantial disagreement about whether the shortfall in participation reflected cyclical factors, which could later rebound, or structural factors, which would keep the LFPR depressed ([Council of Economic Advisers, 2014](#); [Aaronson et al., 2014b](#); [Erceg and Levin, 2014](#); [Krueger, 2017](#)).

The actual path of the LFPR following the Great Recession lines up closely with our estimates, implying that much of this shortfall reflected cyclical factors. We scale our estimates to a Great-Recession-sized shock and compare them to the national age-adjusted LFPR from 2007 through 2019.² The two are remarkably close: the predicted LFPR declines slowly yet persistently through 2014 and then rebounds over the subsequent several years, the same as the actual path. By the end of 2019, only a small portion of the national age-adjusted LFPR is left unexplained by our cyclical model, implying that the response of the LFPR after the Great Recession largely did not reflect any unusual features of this recession and instead was in line with the typical business cycle pattern.

Why does the LFPR typically take so long to recover? We distinguish between two possible explanations. The first, which we term “shadow unemployment”, refers to nonparticipants who are effectively unemployed. The distinction between unemployment and nonparticipation is notoriously subjective, and many people

¹This estimated coefficient on the unemployment rate is at the low end of the range of Okun’s law coefficients estimated in the literature ([Ball, Leigh and Loungani, 2017](#)), supporting that our method measures cyclicalities accurately.

²As we discuss in [Section 4](#), the national age-adjusted LFPR is the correct benchmark to compare our estimates to, and this controls for the aging of the baby boom generation over this period.

may want work in any given month even though they are not currently searching, and then can end up misclassified as nonparticipating (Abowd and Zellner, 1985; Elsby, Hobijn and Şahin, 2015). The second explanation we term “persistent non-market-work activities”, which includes individuals enrolled in school, at home taking care of family, or other similar activities. Although, many of these individuals transition into employment in any given month, those transitions may not respond quickly to changes in labor market conditions, since these activities may take time to enter or exit.

We find that changes in shadow unemployment do not explain the delayed recovery of the LFPR. Although shocks do lead to increases in nonparticipants who self-report that they “want a job”, this type of nonparticipation returns to its pre-shock level at the same time as the unemployment rate and well before the LFPR has fully recovered. Similarly, we also find that shocks lead to an increase in churn between unemployment and nonparticipation, which we take as a measure of shadow unemployment, but this too subsides before the LFPR is fully recovered.

Instead, the delayed cyclical recovery is driven by persistent non-market-work activities, which build in response to a shock but take some time to unwind. Initially following a negative shock, the increase in persistent non-market-work activities is mainly driven by people either taking care of the home and family or going to school. These increases only start to dissipate several years after the shock, reflecting the stickiness of these choices to leave the labor force once they are made. Only after the labor market is well on its way to recovery do these types of nonparticipation return to pre-shock levels. This is also consistent with the patterns we document for flows from nonparticipation to employment: these flows drop initially in response to the shock and then surge only after the unemployment rate has fully recovered, driving the delayed recovery of the LFPR.

Our main approach controls for changes in the composition of state-level populations through age adjustment. This approach removes any mechanical effect on the LFPR from changes in the age structure of the population following local-level shocks, which may occur due to either in-migration or out-migration among particular age groups. In our baseline specification, we adjust for age by residualizing the individual-level LFPR on single-year-age-by-sex fixed effects; we use the state-year average of these age-adjusted outcomes as the dependent variable in our local projections regressions. In this way, our estimates isolate the true cyclical response of the LFPR to an output shock without the influence of compositional changes.

This age-adjustment is necessary, since we show that shocks lead to structural changes in the population of high-LFPR ages. Following a negative 1 percentage

point shock to output growth, the population of 25 to 39 year olds gradually decreases over 10 years, eventually falling up to 4 percent below the pre-shock level, while other age groups see slight increases in population over the same period. Since 25 to 39 year olds tend to have higher LFPRs than other age groups, this response mechanically lowers the unadjusted state-level LFPR in the long run by about 0.2 percentage point in the long run. We explore whether the population changes along other demographic dimensions—such as educational attainment, race, and marital status—but find little further effects beyond age. Our finding that the local LFPR is persistently altered by changes in the population among young, prime-age people adds to the understanding of the migratory adjustment mechanism of local shocks documented by [Blanchard and Katz \(1992\)](#), [Dao, Furceri and Loungani \(2017\)](#), and [Amior and Manning \(2018\)](#).

We also document that the long-lived cyclical nature of the LFPR is especially pronounced for less-advantaged groups in the labor market. Younger workers (ages 16 to 24) exhibit a much larger cyclical response of the LFPR than do prime age workers (ages 25 to 54), while older workers (ages 55+) show a lower degree of cyclical nature. Our estimates show a sharp difference by education level with less-educated prime-age workers experiencing a large decrease in LFPR after a shock, while more-educated workers experience no significant change in labor force participation. We also find substantial inequality in long-lived cyclical nature across racial and ethnic groups, with the prime-age Black LFPR exhibiting larger, longer-lived cyclical nature than the prime-age white LFPR.

Our paper contributes to the literature studying LFPR cyclical nature. Several recent papers take a national-level approach, estimating a structural trend for the LFPR and using deviations from this trend to estimate LFPR cyclical nature ([Aaronson et al., 2014b,a](#); [Council of Economic Advisers, 2014](#); [Krueger, 2017](#); [Montes, 2018](#); [Hornstein and Kudlyak, 2019](#)). This approach requires specifying the structural supply and demand forces that affect participation decisions (see the reviews by [Abraham and Kearney \(2020\)](#) and [Juhn and Potter \(2006\)](#) for a discussion of these forces). Another approach, used by [Aaronson et al. \(2014b\)](#); [Erceg and Levin \(2014\)](#); [Balakrishnan et al. \(2015\)](#), it to rely on state-level variation as we do in our analysis. While some of 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 causally estimating long lags in LFPR cyclical nature. More precisely, unlike the previous papers in this literature, we estimate the dynamic response of LFPR to output shocks by using local projections, which allow for the possibility of very persistent effects on LFPR. Moreover, by using a shift-share instrumental vari-

able approach, we are able to establish a link between exogenous shocks and the dynamic response of LFPR.³ Our approach of using state-level variation to identify LFPR cyclicalities and aggregate up to an estimate of national LFPR cyclicalities follows the growing literature using regional variation to study macroeconomic phenomena (Nakamura and Steinsson, 2014, 2018; Fukui, Nakamura and Steinsson, 2018; Beraja, Hurst and Ospina, 2019; Chodorow-Reich, 2019).

Additionally, following the early work of Blanchard and Katz (1992), several papers investigate how employment adjusts in response to economic shocks at the local level (Decressin and Fatas, 1995; Bound and Holzer, 2000; Dao, Furceri and Loungani, 2017; Amior and Manning, 2018; Hornbeck and Moretti, 2018; Weinstein, 2018; Yagan, 2019; Hershbein and Stuart, 2020) as well as the relationship between shocks and migration (Cadena and Kovak, 2016; Monras, 2018; Howard, 2020). This literature has documented that local labor markets adjust following shocks through changes in migration that return the labor market to equilibrium. We contribute to this literature by showing that this migration channel can have persistent effects on LFPR through altering the composition of the population, primarily among 25 to 39 year olds, which makes it important when studying local shocks to use the age-adjusted LFPR.

We also contribute to a literature following Okun (1973) that examines how different demographic groups fare during a long recovery. Some groups of workers may disproportionately benefit from a tight labor market, as discussed in Bradbury (2000); Hoynes (2000); Jefferson (2008); Hoynes, Miller and Schaller (2012); Wilson (2015); Cajner et al. (2017); Hotchkiss and Moore (2018); Fallick and Krolikowski (2019) and Aaronson et al. (2019). Some of these differences in the benefits of a tight labor market may stem from delayed recoveries of LFPR, since we estimate that some groups' LFPRs take substantially longer to recover after a typical recession. These differences in LFPR cyclicalities make it necessary to sustain recoveries beyond the point at which unemployment has fully recovered if policymakers want to ensure all groups experience a full recovery.

2 Research Design

We measure the cyclicalities of labor force participation by estimating its response to state-level business cycles in order to sidestep the issue of trend changes in participation, which complicate identifying cyclicalities at the national level. In this section,

³Balakrishnan et al. (2015) also use a similar shift-share instruments, but for employment, while our paper uses output. Using the latter has several methodological advantages as we argue later on.

we outline our research design, starting with the identification problem and our approach to solve 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 assume 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 the LFPR, given the sharp and changing nature of LFPR 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 regressions indexed by k :

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

where $y_{s,t}$ represents the labor market dependent variable of interest—for example, the LFPR—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 $W_{s,t}$ represents a vector of control variables. In our baseline specification, the controls include state and year fixed effects.

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. Nationwide phenomena that affect labor market outcomes across all states equally, including demographic shifts (such as the aging of the baby boom generation) and national policy responses (such as monetary policy shocks), are controlled for nonparametrically by this approach.

We view local projections regressions as a better alternative in our setting than vector autoregressions (VARs). [Stock and Watson \(2018\)](#) point out that instrumented versions of VARs and local projections identify the same IRFs under standard conditions, but VARs may not correctly identify IRFs if the true IRFs are not invertible. In terms of efficiency, instrumented local projections have the same properties as VAR models with internal instruments, as documented by [Plagborg-Møller and](#)

Wolf (2021). Additionally, Olea and Plagborg-Møller (2020) show that local projections have attractive properties for inference.⁴ For these reasons, we use local projections in our main specification, but we return to the question of whether VARs are appropriate for our setting in Section 9.4, where we conduct a test of invertibility from our estimated IRFs.

Shocks: We measure the business cycle shock using real gross state product (GSP) growth as estimated by the BEA. Specifically, we define $\text{Shock}_{s,t} \equiv \Delta \text{GSP}_{s,t}$, where $\Delta \text{GSP}_{s,t}$ is the year-over-year percent change in GSP. That is, the shock is a one-time, temporary, one percentage point shock to GSP growth. All else equal, the shock leads to a permanently lower level of output.⁵

GSP estimates are based on the factor incomes earned and other costs incurred in production, which is the same concept for measuring output as is used by Gross Domestic Income (GDI) at the national level. For each state, GSP sums labor income, capital income, and business taxes, where each of the three components is estimated by industry. Note that labor income is based on wage and salary accruals (as opposed to disbursements), which implies that retroactive wage payments (bonuses) are counted for the year in which they were earned rather than when they were received.

We view our choice to define business cycles based on output as superior to alternative approaches that use employment. Using GSP provides a measure of business cycle fluctuations at the state level that is more comprehensive than only using employment, which omits fluctuations in productivity. Additionally, if shocks take time to propagate to the labor market, using output will correctly time business cycles, while employment-based business cycles will tend to lag behind the true timing of the shock. Lastly, estimating the response of LFPR to an output shock, rather than an employment shock, makes the results more interpretable in the context of Okun’s Law, a key economic relationship used among many policymakers.

Potential Endogeneity: The coefficient $\beta^{(k)}$ gives the k -period-later response of y to a one-time, temporary, one percentage point shock to GSP growth.⁶ For $\beta^{(k)}$ to

⁴Herbst and Johannsen (2020) show that local projections can be biased in small samples when the outcome variable is highly persistent, but this bias is likely to be minimal in our setting. The dependent variable in our regressions, the state-level age-adjusted LFPR, only has an autocorrelation coefficient of 0.81, lower than the 0.9–0.99 range in which this bias becomes acute.

⁵We show in Section 8 that the lower level of output in large part reflects a permanently lower level of output per employee.

⁶Assuming that the shock leaves GSP growth in other periods unaffected, this results in a permanent one percent shock to the level of GSP.

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:

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

However, two key concerns suggest this requirement might not be met in practice. One concern is that employment may affect GSP, as lower employment (through higher unemployment, lower LFPRs, or both) will lower GSP if productivity is held constant. 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 to isolate shocks to demand at the state-level. The first-stage equation is as follows,

$$\Delta GSP_{s,t} = \alpha Z_{s,t} + \gamma W_{s,t} + \nu_{s,t} \quad (2)$$

where

$$Z_{s,t} \equiv \sum_q \Delta GDI_{q,-s,t} \omega_{q,s,t-5}. \quad (3)$$

Industries are indexed by q , and $\omega_{q,s,t-5}$ represents the three-year moving average of industry q 's share of total GSP in state s five years previously.⁷ $\Delta GDI_{q,-s,t}$ represents the growth rate of national gross domestic income in industry q for period t using the "leave-one-out" approach—that is, we calculate $GDI_{q,-s,t}$ by summing up $GSP_{q,s,t}$ across all states except for state s .

This formulation of the shift-share instrument relies on industry variation in output, rather than employment. Many previous studies, including [Blanchard and Katz \(1992\)](#), [Dao, Furceri and Loungani \(2017\)](#), [Adão, Kolesár and Morales \(2019\)](#), and [Goldsmith-Pinkham, Sorkin and Swift \(2020\)](#), measure the response of employment to a shift-share instrument that uses industry variation in employment. However, we view industry variation in output as more appropriate for our setting, both because changes in output are likely to be more closely aligned to industry cycles and because output is a distinct variable measured separately from our outcomes of interest.

⁷During the first five years of available industry data, we calculate $\omega_{q,s,t-5}$ from industry q 's share of total GSP in the first year of data instead.

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

$$\mathbb{E}[Z_{s,t} \cdot \Delta GSP_{s,t} | W_{s,t}] = \alpha \neq 0 \quad (\text{relevance}) \quad (4)$$

$$\mathbb{E}[Z_{s,t} \cdot \epsilon_{s,t} | W_{s,t}] = 0 \quad (\text{contemporaneous exogeneity}) \quad (5)$$

$$\left. \begin{aligned} \mathbb{E}[Z_{s,t} \cdot \epsilon_{s,t+k} | W_{s,t}] &= 0 \\ \mathbb{E}[Z_{s,t} \cdot \Delta GSP_{s,t+k} | W_{s,t}] &= 0 \end{aligned} \right\} \text{for } k \neq 0 \quad (\text{lead-lag exogeneity}) \quad (6)$$

$Z_{s,t}$ 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 meet the relevance condition since local output in a given industry is likely to be correlated with national output in that industry due to changes in industry technology or relative demand. The contemporaneous exogeneity assumption will hold as long as the national industry shocks used to construct $Z_{s,t}$ 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 $Z_{s,t}$ 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. In Section 8 and Section 9, we show that $\mathbb{E}[Z_{s,t} \cdot \Delta GSP_{s,t+k} | W_{s,t}] \approx 0$ for $k \neq 0$ in our sample, confirming this aspect of lead-lag exogeneity.⁸

There are multiple interpretations of exogeneity for the shift-share instrument. The variation in the shift-share instrument comes from differential exposure to national shocks across regions based on initial industry shares. Goldsmith-Pinkham, Sorkin and Swift (2020) point out that this variation is equivalent to instrumenting with the industry shares directly, and therefore exogeneity of the instrument comes from exogeneity of these shares. Borusyak, Hull and Jaravel (forthcoming) provide an alternative interpretation in which the national shocks are exogenous.

2.2 Inference

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

⁸If the shock were positively autocorrelated, then some of the effect estimated by $\beta^{(k)}$ would be the result of the persistence of the shock, thus biasing our estimate upward. Conversely, if the shock were negatively autocorrelated, the effects of an output shock on the LFPR would be stronger than what is estimated by $\beta^{(k)}$. Since we find that autocorrelation in $Z_{s,t}$ is minimal, this bias doesn't affect our estimates.

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 \(2019\)](#) raise concerns that this approach may understate uncertainty in shift-share designs. However, these concerns are primarily about settings where variation comes from a subset of industries, while our setting uses the full set of industries.⁹ We validate this choice in [Section 9.3](#) with a placebo exercise, which indicates that our clustered standard errors are, if anything, a bit conservative for this setting.

Weighting: We weight each regression of outcome $y_{s,t}$ for group j by the population, n_{st}^j 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 LFPR. Second, the smallest states have relatively few respondents in the CPS, which has the potential to generate noise when calculating state-level LFPRs for those smaller states and yield imprecise regression estimates. 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 influence of noise in our estimates.

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 \(1997\)](#). Although the instrument and endogenous variable are the same in all specifications, the F-statistics may vary across regressions for different demographic groups due to the different state population weights for different groups.

2.3 Data

We combine state-level data from multiple sources to form an annual panel. Labor market outcome variables consist of the unemployment rate, the labor force participation rate, and the employment-to-population ratio, each of which is calculated from Current Population Survey microdata. For each rate, we compute the average

⁹Our analysis is most similar to the results shown in Panel B of Table 6 in [Adão, Kolesár and Morales \(2019\)](#), which shows that more sophisticated approaches to estimate confidence intervals are not meaningfully different from clustering by local labor market.

over the calendar year in each state. Our main specification uses the CPS sample of civilian noninstitutionalized people ages 16 and over to compute each of these rates. In later sections, we compute these rates for subgroups of the population.

In order to control for shifting demographics unrelated to the business cycle, we age-sex-adjust each of our labor market outcome variables. That is, for an outcome $y_{i,s,t}$ for person i in state s and year t , we estimate the the following equation on our CPS sample:

$$y_{i,s,t} = \theta_{\text{age}(i),\text{sex}(i)} + \tilde{y}_{i,s,t} \quad (7)$$

where $\theta_{\text{age}(i),\text{sex}(i)}$ is a age-by-sex fixed effect. We then compute the average age-adjusted outcome for state s in year t as

$$\tilde{y}_{s,t} \equiv \sum_{i \in (s,t)} \tilde{y}_{i,s,t} w_{i,s,t} \quad (8)$$

where $w_{i,s,t}$ is the CPS sampling weight for person i . This procedure removes changes from our outcomes that are due to changes in the age structure of the population such as the aging of the baby boom generation, which has been shown to be responsible for variation in labor market outcomes over time (see, e.g., [Shimer, 1999](#)). We use the age-adjusted rates in all of our main estimates, but return to examine the role of this adjustment compared to alternative adjustments and unadjusted rates in [Section 6.1](#).

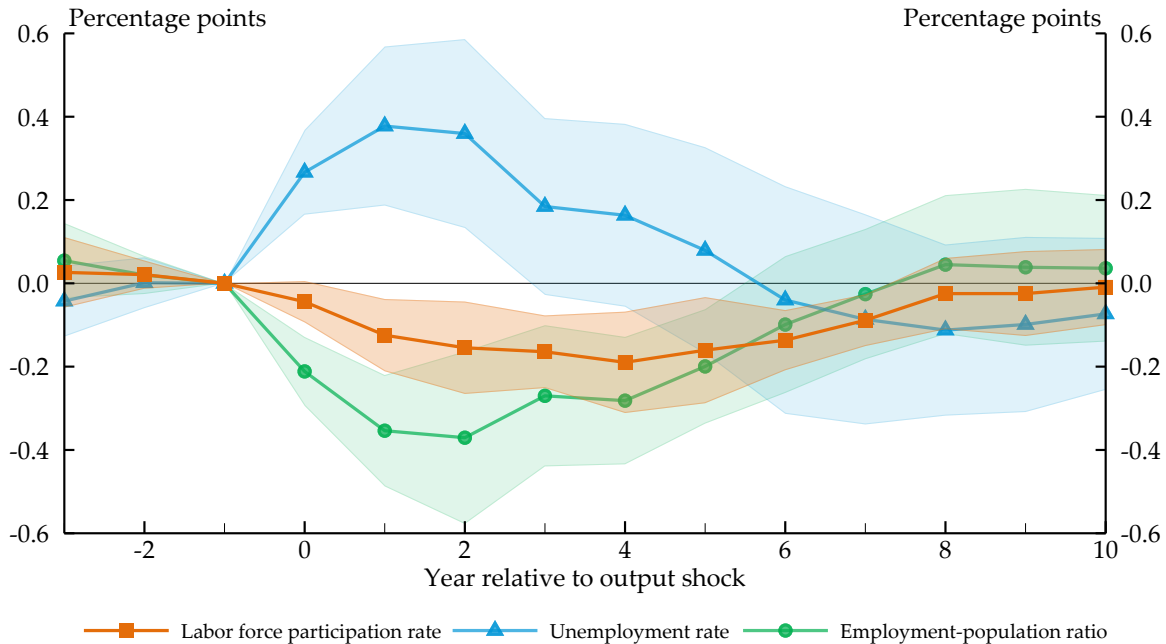
Data on GSP for each state and year are obtained from the BEA. GSP data by industry are from the BEA as well, using SIC-coded industries for 1976–1998 and NAICS-coded industries for 1998–2018. For the purposes of decomposing the variation in the shift-share instrument in [Section 8.3](#), we link a subset of industries between SIC and NAICS that are categorized in essentially the same way in both systems, and treat all other industries as distinct between the two systems.

3 Cyclicity of Labor Market Outcomes

[Figure 1](#) presents our estimates of the impulse response functions for the age-adjusted LFPR, unemployment rate, and employment-to-population ratio (EPOP) from 3 years before the shock to 10 years after the shock. For ease of interpretation, we report all of our estimates as the response to a temporary negative 1 percentage point shock to GSP growth, so that the cyclical responses will have the same sign as in a recession.

The unemployment rate, LFPR, and EPOP all respond to cyclical shocks, but with varying timing. For the unemployment rate, a contractionary 1 percentage

Figure 1: Estimated Cyclical Responses to a Negative Output Shock



Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GDP growth in year 0. All outcomes are adjusted for changes in the age-by-sex composition of the population. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, and authors' calculations.

point shock to output growth causes a contemporaneous 0.25 percentage point increase in the unemployment rate. The increase in the unemployment rate continues in the following year and peaks at 0.4 percentage point one year after the shock. Our estimate of the total increase in the unemployment rate due to a negative 1 percentage point shock to GDP is at the low end of the range of Okun's law coefficients estimated in the literature of -0.5 to -0.4; see, for example, Ball, Leigh and Loungani (2017). Following the peak one year after the shock, the unemployment rate steadily declines by about 0.1 percentage point per year until it returns to its pre-shock value about six years after the shock and remains there. This asymmetric response of a sharp increase followed by a gradual decrease is consistent with the "plucking" dynamics of business cycles examined by Dupraz, Nakamura and Steinsson (2019).

The LFPR also shows a significant response to a negative shock, but with a substantial delay compared to the unemployment rate. Specifically, the LFPR declines by less than 0.1 percentage point in the year of the shock, much smaller than the increase in the unemployment rate. However, while the unemployment rate quickly

peaks and begins to recover, the LFPR continues to steadily decline for several years after the shock, finally reaching a trough four years later at a level that is 0.2 percentage point below its initial value. After reaching its trough, the LFPR gradually recovers and only attains its pre-shock level eight years after the initial shock, two years after the unemployment rate has fully recovered.

The different patterns for the LFPR and unemployment rate reflect different *cyclical profiles*, which we can show formally with a nonlinear Wald-type test. We denote the set of coefficients tracing out the impulse response of the LFPR as $\{\beta_{\text{LFPR}}^{(k)}\}$ and the set of coefficients for the unemployment rate as $\{\beta_{\text{UR}}^{(k)}\}$. Our null hypothesis is that the LFPR response has the same time profile as the unemployment rate but perhaps a different cyclical loading: $\beta_{\text{LFPR}}^{(k)} \equiv \frac{\beta_{\text{UR}}^{(k)}}{\phi}$ for each horizon k . Under this null, the ratio of coefficients $\frac{\beta_{\text{UR}}^{(k)}}{\beta_{\text{LFPR}}^{(k)}}$ is the same at every horizon k . To test this, we stack the samples used to estimated impulse responses for both variables and re-estimate [Equation 1](#), from which we obtain a covariance matrix containing all coefficients for both impulse responses.¹⁰ We use the delta method to construct a nonlinear Wald-type test statistic for the restriction that the ratio of coefficients is the same at each horizon. For the null hypothesis that lags 1 to 8 share the same ratio, we obtain a test statistic of 31.69 with a p-value of 0.000, enough to strongly reject the null hypothesis that the time profile is the same for both variables.

The combination of the LFPR and unemployment rate responses create cyclical-ity in the EPOP that is both large and long-lasting. The EPOP declines rapidly at the onset of the shock, reflecting the initial spike in the unemployment rate, and reaches its trough at about -0.4 percentage point two years after the shock. Thereafter, the EPOP steadily recovers by about 5 basis points per year until it is fully recovered seven years after the shock. While the EPOP shortfall in earlier years reflects high unemployment, the remaining EPOP shortfall in years 5 to 7 is almost entirely accounted for by the LFPR.

4 Implications for National LFPR Cyclicity

In this section we lay out a framework for aggregating our results from state-level business cycles to the national level. Using this framework, we show that our estimates closely match the observed dynamics of the LFPR following the Great Recession. In particular, the strength in labor force participation starting in 2014 lines up

¹⁰This is equivalent to a seemingly unrelated regression with the same right-hand-side variables in each equation ([Davidson and MacKinnon, 1993](#)).

closely with the delayed recovery of the LFPR in our estimates.

The aggregate LFPR at the state level is an average across the LFPRs of all ages, a , within the population. The LFPR for each of these ages in turn can be broken down into a cyclical component and a secular trend component unrelated to the business cycle. Let $C_{s,t}$ denote the measure of the business cycle in state s and time period t , let $\beta(L)$ denote the cyclical coefficients tracing out the impulse response, and let $\alpha_{a,s}$ and $\alpha_{a,t}$ be state-age fixed effects and year-age-specific secular trend components respectively. With a set of population weights for each age $w_{a,s,t}$, the state-level LFPR can be written as:

$$\text{LFPR}_{s,t} = \sum_a (\alpha_{a,s} + \alpha_{a,t} + \beta(L)C_{s,t}) w_{a,s,t}$$

From this expression, the first-order approximation for changes in the state-level LFPR from period t to $t+k$ can be broken down into three components¹¹:

$$\Delta \text{LFPR}_{s,t+k,t} \approx \underbrace{\sum_a (\Delta \alpha_{a,t+k,t} w_{a,s,t})}_{\text{Trend}} + \underbrace{\beta(L) \Delta C_{s,t+k,t}}_{\text{Cycle}} + \underbrace{\sum_a (\alpha_{a,s} + \alpha_{a,t} + \beta(L)C_t) \Delta w_{a,s,t+k,t}}_{\text{Population changes}}$$

When using our methodology to estimate the cyclical response of the LFPR, the first and third terms drop out. The first term consists of national trends in LFPR, which can be broken down into a purely national component $\sum_a (\Delta \alpha_{a,t+k,t} w_{a,t})$ that is absorbed by our time fixed effects and a residual component $\sum_a (\Delta \alpha_{a,t+k,t} (w_{a,s,t} - w_{a,t}))$. Our identification assumption implies that the residual component is equal to zero.¹² The third term in the decomposition above is equal to zero in our setting since we use the age-adjusted LFPR as the outcome, which mechanically adjusts for $\Delta w_{a,s,t+k,t}$.¹³

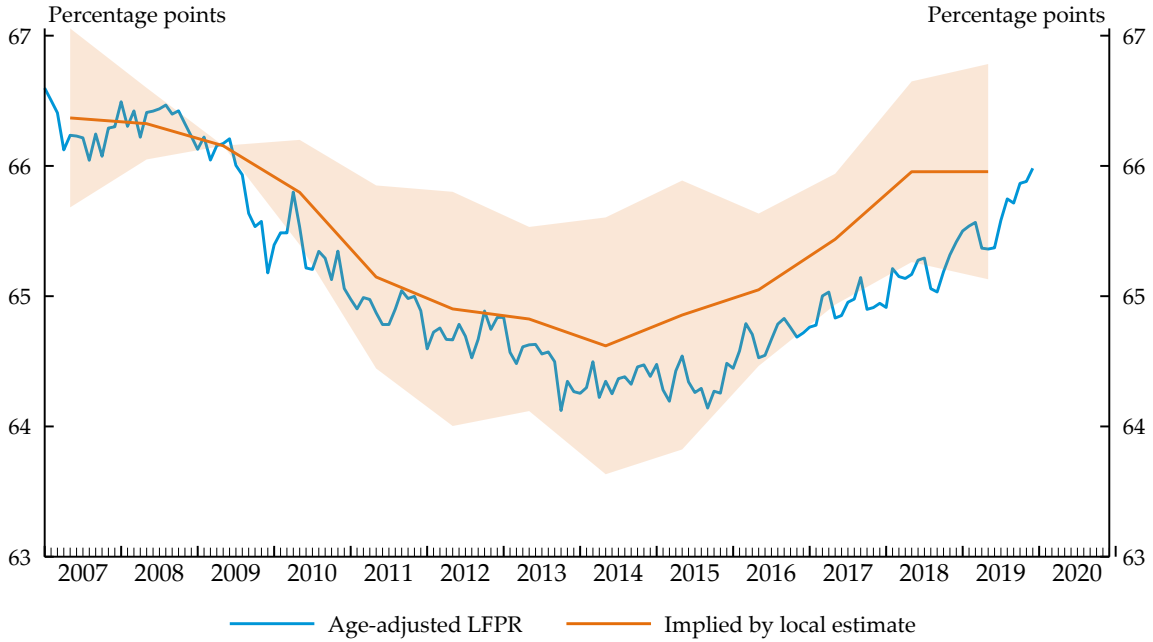
As a result, our methodology provides estimates for the cyclical coefficients $\hat{\beta}(L)$. These coefficients represent the lagged change in LFPR associated with a one unit change in output. Importantly, since we include time fixed effects, the change in output $\Delta C_{s,t+k,t}$ is measured relative to nationwide trend output growth, which accounts for gradual increases in GDP due to population growth and productivity, among other forces.

¹¹The approximation omits a higher-order term involving the product $\Delta \alpha_{a,t+k,t} \cdot \Delta w_{a,s,t+k,t}$.

¹²There is no particular reason why the residual component would not be equal to zero, and this assumption is consistent with the lack of noticeable pre-trends in our LFPR estimates.

¹³Note that even if we had not used the age-adjusted LFPR, our inclusion of time period fixed effects would absorb any of the variation in $\Delta w_{a,s,t+k,t}$ that is common nationwide, for example the aging of the baby boom population. However, $\Delta w_{a,s,t+k,t}$ also includes age-selected migratory responses, which are *not* controlled for by time period fixed effects, but are controlled for by using the age-adjusted LFPR.

Figure 2: Actual and Predicted LFPR after the Great Recession



Note: The blue line shows the LFPR adjusted for changes in the age composition of the population since 2007. The orange line uses our main estimates of Equation 1 scaled by the -8.1% decline in GDP relative to its estimated trend during the Great Recession (CBO, 2007). The bands around this line show a 95% confidence interval, based on standard errors clustered by state.

Source: BLS, BEA, and authors' calculations.

The national LFPR is an average of state LFPRs, and can be decomposed similarly. Averaging over states and dropping the state subscripts, the decomposition above becomes:

$$\Delta \text{LFPR}_{t+k,t} \approx \underbrace{\sum_a (\Delta \alpha_{a,t+k,t} w_{a,t})}_{\text{Trend}} + \underbrace{\beta(L) \Delta C_{t+k,t}}_{\text{Cycle}} + \underbrace{\sum_a (\alpha_{a,t} + \beta(L) C_t) \Delta w_{a,t+k,t}}_{\text{Population changes}}$$

From this decomposition, we use our estimated coefficients $\hat{\beta}(L)$ to trace out the predicted change in the national age-adjusted LFPR to a recessionary shock, namely the Great Recession. Using the age-adjusted LFPR removes the third term, and we pick a time period (2007–2019) over which trends within age groups are estimated to have been close to zero on net (Montes, 2018), removing the first term. Additionally, by using the age-adjusted LFPR, we remove any contribution of migratory responses to shocks on the LFPR. While we find significant migratory responses affecting the LFPR at the state level, as we will discuss in detail in Section 6.2, those migratory responses are unlikely to have an effect on the national LFPR to the extent

that this reflects only interstate migration and not international migration.

Figure 2 plots our main estimates of the cyclical response of LFPR applied to the Great Recession shock along with the actual age-adjusted LFPR for this time period. For the Great Recession shock, we compute the decline from 2007:Q4 to 2009:Q2 in real GDP from BEA, minus the expected increase in real potential output over the same period from CBO’s August 2007 projections (CBO, 2007).¹⁴ We use the change relative to trend in order to align with our estimates, which control for trends in potential output through year fixed effects. The predicted path of the LFPR from our estimates in Figure 1 is $-8.1 \times \hat{\beta}^{(k)}$ for each horizon k , and we plot this against the actual path of the LFPR in Figure 2.

The prediction from our estimates is remarkably close to the actual age-adjusted LFPR over this period, featuring a similar slow decline over 2009–2014 and subsequent rebound in later years. By 2019, only a small portion of the LFPR recovery is left unexplained by our model. This similarity suggests that the LFPR largely followed its usual cyclical dynamics over this period with little deviation.

5 What Drives the Long-Lived Cyclicity of Labor Force Participation?

In this section, we draw on two additional features of the CPS data to understand why the LFPR exhibits long-lived cyclicity. First, we use individuals’ self-reported reasons for nonparticipation to examine the cyclicity of “discouraged” individuals—those who report wanting a job but remain out of the labor force—compared to individuals engaged in non-market-work activities, such as home production or schooling. Second, we examine the cyclicity of flows into and out of the labor force to understand whether the shortfall of participation is caused more by a lack of individuals joining the labor force or a surplus of individuals leaving the labor force.

5.1 Reasons for Labor Force Nonparticipation

Business cycle shocks may lead people to make decisions that have persistent effects on their labor supply, which could account for the long-lived cyclicity in the LFPR. Such decisions may include enrolling in school, staying at home and taking care for a family member, applying for disability benefits, or retiring. Alternatively,

¹⁴We use the August 2007 projections in order to ensure that the potential output estimates are not affected by the downturn in LFPR that occurred after the recession.

the long-lived cyclical in the LFPR may reflect individuals becoming discouraged and stopping their search for work, even though they would still prefer to be employed.

To determine the extent to which each of these explanations may account for long-lived cyclical, we use questions in the CPS that ask nonparticipants about their reason for being out of the labor force. Throughout the sample period, nonparticipants were asked whether they want a job, which provides an indication of desired labor supply. Additionally, from 1989 onward nonparticipants were asked to categorize their main reason for being out of the labor force between being ill or disabled, in school, taking care of home or family, retired, or other, and this question is a full partition of the not-in-the-labor-force group.¹⁵ For each of these questions, we compute the share of the population in each state-year that is made up by nonparticipants in each category, and estimate Equation 1 using these outcomes. The estimated impulse responses are shown in Figure 3. We show the IRFs only through eight years following the shock, since the estimates around lag eight become extremely noisy due to the limited sample.

Increases in schooling, staying at home due to family responsibilities, and rising self-reported disability all play important roles in shaping the cyclical response of aggregate labor force participation.¹⁶ Initially, nonparticipants taking care of home/family constitute the largest response, with schooling close behind. However, nonparticipants reporting illness or disability grow steadily in response from year two onward, and comprise a larger portion of the response in years 5 to 7 than people taking care of home or family. People in school grow steadily as well, before falling rapidly in years 7 and 8 when the overall LFPR is reaching its pre-shock level.

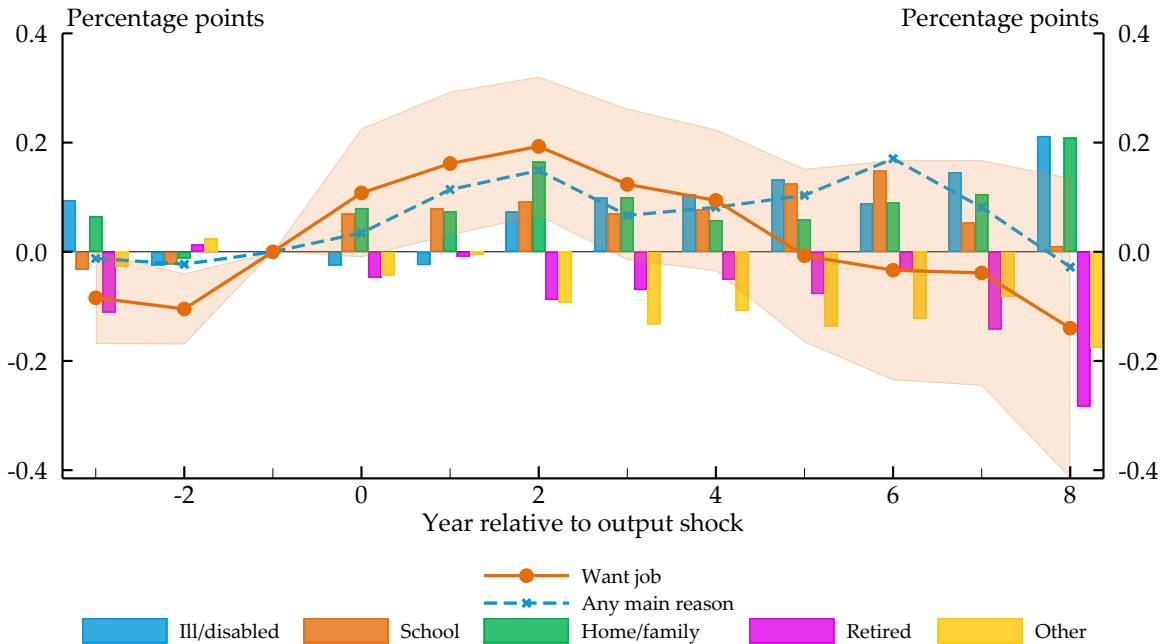
Interestingly, the cyclical response of labor force participation does not seem to be driven by retirement decisions. If anything, retirements appear to exert an upward pressure on the LFPR. This could indicate that recessions induce individuals to postpone retirements, perhaps due to a fall in the value of their retirement savings or to potentially offset income losses of their household members who may lose a job.

Separately, we also look at the cyclical of labor force nonparticipants who say

¹⁵For both of these questions, surveys before 1994 only asked these questions to the roughly $\frac{1}{4}$ of nonparticipants who are part of the Outgoing Rotation Groups in months 4 and 8 in sample. Further, the “want a job” question is separate from the “main reason for being out of the labor force” question (e.g. some respondents who report being in school may also report wanting a job, while others in school may report not wanting a job).

¹⁶In unreported results we find that increases in schooling are most prominent for young people, but also notable for prime-age individuals. Rising disability is mostly present for individuals aged 55 years and over.

Figure 3: Cyclicity by Self-Reported Reason for Labor Force Nonparticipation



Note: Each line and set of bars shows the estimated coefficients from Equation 1 using as the outcome the share of the population out of the labor force and reporting the specified reason. The band around the orange solid line shows a 95% confidence interval, based on standard errors clustered by state. Reporting “want job” is not exclusive with reporting any of the main reasons. The blue dashed line is equal to the sum of the bars in each period. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. All outcomes are adjusted for changes in the age-by-sex composition of the population. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population. Standard errors clustered by state. *Source:* BLS, BEA, and authors’ calculations.

they want a job, which can represent labor market slack. Although nonparticipants who want a job drive essentially all of the early rise in nonparticipation, their participation recovers faster than nonparticipation as a whole, reaching its pre-shock level around the same time as the overall unemployment rate does (years 4–5). This suggests that expansive definitions of the unemployment rate that include nonparticipants who want a job—BLS’ U-5 measure includes some of them—are able to capture additional cyclicity beyond the main unemployment rate, but do not capture the long-lived cyclicity of participation.

5.2 Labor Market Flows

Additionally, the panel structure of the CPS allows us to examine the contributions of inflows and outflows to the long-lived cyclicity of the LFPR. Examining the flows provides more insight into the cyclicity of stocks, as demonstrated in pre-

vious work including [Elsby, Hobijn and Şahin \(2015\)](#), [Elsby et al. \(2019\)](#), and [Cairó, Fujita and Morales-Jiménez \(2021\)](#). We calculate annual labor market transitions by matching individuals in the CPS over 12-month horizons, and express those flows as shares by dividing by population 16 years and over. To be consistent with our baseline results, we also adjust for age by residualizing the flow rates using person-level data to net out the composition component explained by the age distribution of the people in each state.

Three aspects of the response of flows (shown in [Figure 4](#)) are worth noting.¹⁷ First, at the onset of a negative output shock, labor force entry drops, driven by a large decline in the flow from nonparticipation to employment, which could reflect decisions to prolong schooling or to stay home and take care of family, as discussed in the previous subsection. Second, from years 1 to 2 after the shock, flows between unemployment and nonparticipation in both directions rise notably. In particular, negative business cycle shocks lead to an increase of “in-and-outs”, that is individuals who temporary leave the labor force, perhaps due to discouragement.¹⁸ Flows between unemployment and nonparticipation remain elevated until roughly year 6, which is also how long the unemployment rate remains elevated (recall [Figure 1](#)). Finally, flows from nonparticipation to employment eventually surge around 8 years after the shock, which leads to the recovery of the LFPR. In terms of magnitudes, outflows are elevated by 2 to 3 basis points after a negative shocks, while inflows are depressed by about 5 basis points. Cumulatively, the net effect for flows is similar to the one estimated for the stock of labor force participants shown in [Figure 1](#).

6 The Role of Changing Demographic Composition

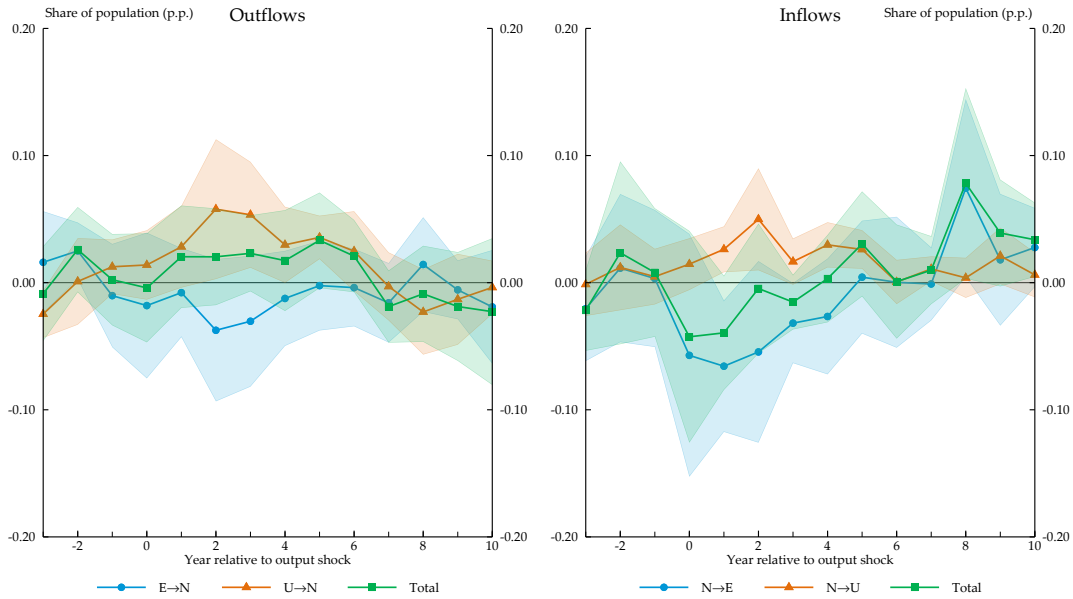
In addition to changes in the age-adjusted LFPR, shocks may lead to changes at the state level in the age structure of the population or other demographics. In this section, we examine how the demographic composition of the state-level population responds to output shocks, finding evidence that shocks induce permanent, structural composition shifts away high-LFPR subgroups in affected states.

We start by showing that the unadjusted LFPR experiences a persistent shortfall

¹⁷For brevity, we do not report flows between employment and unemployment, since these are neutral with respect to the LFPR.

¹⁸Note that while the hazard rate of $U \rightarrow N$ flows ($U \rightarrow N$ flow divided by the stock of unemployment) declines in recessions, in part driven by compositional changes of unemployed and their higher eligibility for unemployment insurance, the stock of unemployed rises even more during recessions, leading to an increase of $U \rightarrow N$ flows as a share of the population ([Elsby, Hobijn and Şahin, 2015](#)).

Figure 4: Estimated Cyclical Responses of Flows to a Negative Output Shock



Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market flow. Flows are measured as the share of the population experiencing the specified type of flow from the beginning of a 12-month period to the end. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. All outcomes are adjusted for changes in the age-by-sex composition of the population. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population. *Source:* BLS, BEA, and authors' calculations.

after output shocks. However, this persistent effect is not the result of hysteresis but instead reflects changes in the demographic composition of the population at the state level, primarily the age distribution. We find little to no contribution from changes in education, race, ethnicity, and marital status.

Next, we examine how the state-level population in each single-year age group changes in response to output shocks, finding that declines are concentrated among 25 to 39 year olds. Since this age group tends to have higher LFPRs than other age groups, declines in its population pull down the unadjusted overall LFPR mechanically after an output shock. We caution that this phenomenon raises the importance of using age-adjusted LFPRs to examine questions about cyclical and hysteresis in response to local shocks.

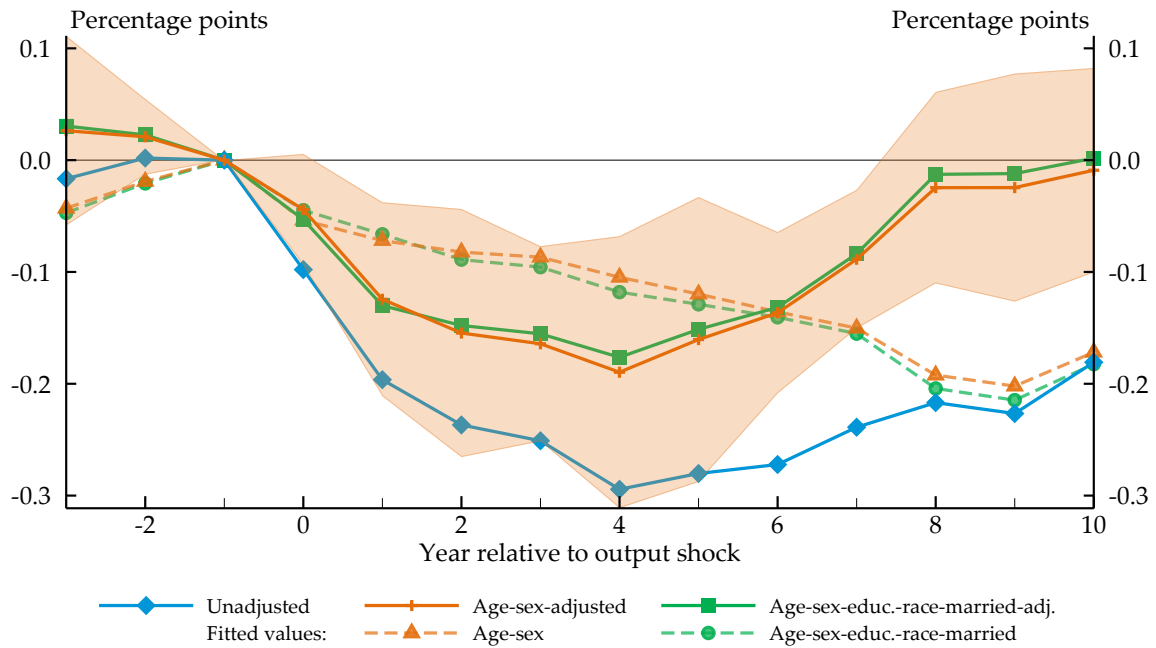
6.1 Cyclical of Adjusted and Unadjusted LFPRs

To investigate how demographics affect the cyclical of the LFPR, we compare our age-adjusted baseline estimates to two alternative benchmarks.

First, we estimate Equation 1 using the unadjusted LFPR. Figure 5 shows that the unadjusted LFPR steadily declines to its trough in year four, with similar timing but a steeper decline compared to the age-adjusted LFPR. However, while the age-adjusted LFPR subsequently recovers back to its pre-shock level, the unadjusted LFPR merely edges up a bit, but remains well below its pre-shock level even ten years after the shock.

While a persistent shortfall of the unadjusted LFPR after a shock might be interpreted as evidence of hysteresis, we caution that this is not the case in our setting. By hysteresis, it is commonly meant that people become persistently less likely to participate in the labor market as a result of the shock. However, our estimates do not suggest that people experience persistently lower participation conditional on their demographics, as the age-adjusted LFPR fully recovers on average by 8 years after a shock.

Figure 5: Cyclicity by Demographic Adjustment



Note: Each line shows the estimated coefficients from Equation 1 using the specified adjusted, unadjusted, and fitted-value LFPR as the outcome. The band around the orange solid line shows a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population. Source: BLS, BEA, own calculations.

For our second benchmark, we consider a broader adjustment for multiple demographic characteristics. Using person-level data from the CPS, we regress a person's labor force participation indicator on demographic characteristics using the

following linear-probability model:

$$Y_{i,s,m,t} = \psi_0 + \Psi^{i,m,t} D_{i,m,t} + \Psi^{s,m,t} W_{s,m,t} + \eta_{i,s,t} \quad (9)$$

where $Y_{i,s,m,t}$ is a indicator variable indicating whether person i in state s was participating in the labor force in month m of year t ; $D_{i,m,t}$ is a vector of indicator variables over the demographic characteristics of person i in month m of year t that include age, gender, educational attainment, race/ethnicity, and marital status; and $W_{s,m,t}$ is a vector of state, month, year fixed effects. The age variables are single-year age indicators for ages 16 to 79 and a indicator variable for ages 80 years and older. The educational attainment indicators partition attainment into five categories: less than a high school degree, a high school degree, some college, a four-year college degree, and more than a college degree. The race/ethnicity indicators partition the population into four groups: non-Hispanic white, non-Hispanic Black, Hispanic, and other. Marital status is a single indicator indicating whether an individual is married. We include month-of-year indicator variables to account for seasonality.

Using the estimated coefficients from [Equation 9](#), we predict whether a person is participating in the labor force based on their demographic characteristics and denote this by $\widehat{Y}_{i,s,m,t}^D$. With this fitted value, we calculate the demographically-adjusted LFPR as the residual, $\widehat{Y}_{i,s,m,t}^{D.adj}$. We then aggregate the person-level fitted and residual components to calculate monthly rates for the fitted and demographically-adjusted labor market variables in each state s , and then average across months within year t to create a fitted value component, $\widehat{y}_{s,t}^D$, and a demographically-adjusted component, $\widehat{y}_{s,t}^{D.adj}$. Finally, we use those fitted values and demographically-adjusted state-level variables as the dependent variable in [Equation 1](#).

The additional demographic controls beyond age make little to no difference in estimating LFPR cyclicalilty. [Figure 5](#) shows that the addition of adjustments for education, race/ethnicity, and marital status results in nearly the same estimated impulse response as our baseline estimates, which adjust for age and sex only. The similarity of adjusted values is mirrored in the fitted values, which both decline steadily in response to the shock. This pattern points to the age structure of the state-level population changing persistently in a way which would mechanically pull down the LFPR absent adjustment.

In [Appendix Figure A.1](#) we repeat this exercise for the unemployment rate. In contrast to the LFPR, we find that demographics explain essentially none of the response of unemployment, both immediately following the shock and in the long-run afterwards.

6.2 Response of Population Composition to Cyclical Shocks

Why does the age-composition of the state-level 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, [Amior and Manning \(2018\)](#) show that long-term adjustment in regional populations tends to differ across demographic groups, and if the migration response to business cycles similarly differs across groups, then the composition of the population could be altered by these shocks. For example, if shocks lead to higher migration responses among prime-age people, who tend to have higher LFPRs, then these shocks could alter the composition of the population resulting in a permanently lower LFPR.

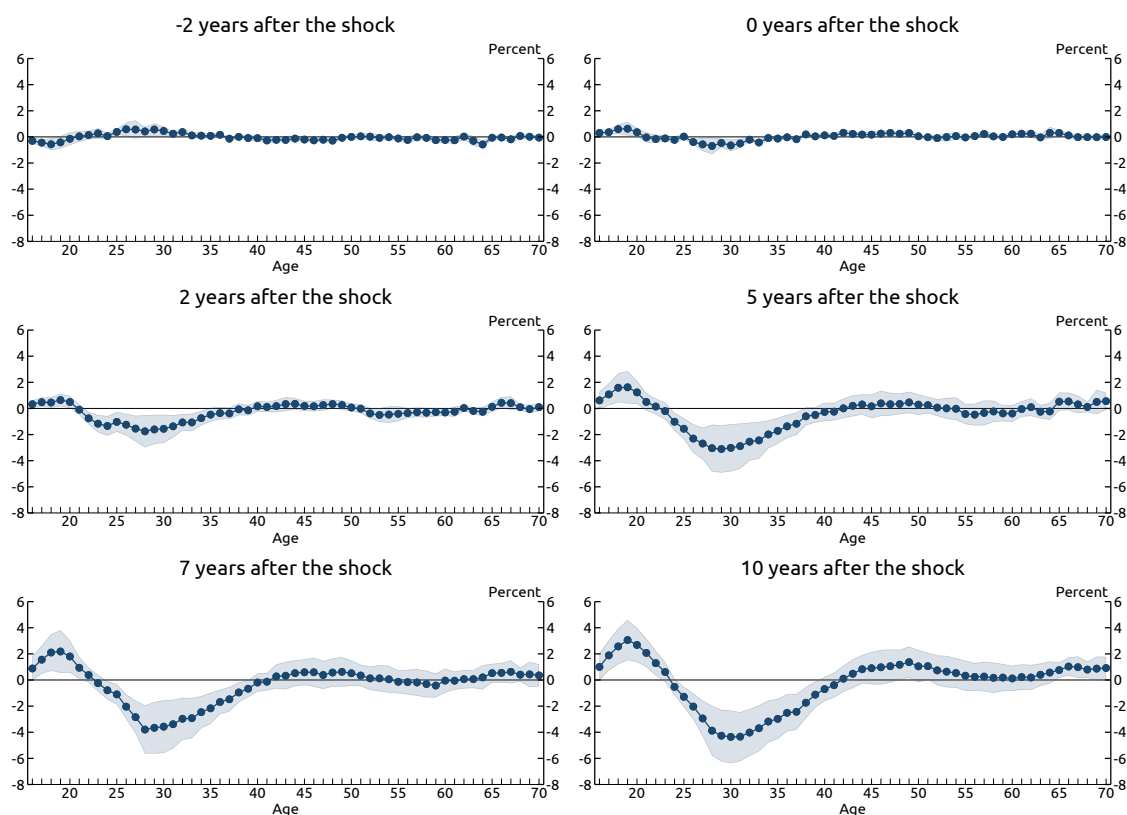
In this section, we examine how the age composition of a state's population across single-year-age groups responds to a business cycle shock. Understanding the changes in the age structure are essential not only for understanding how the population changes but also for understanding how national LFPR cyclicity may be related to local LFPR cyclicity. If shocks induce out-migration of selected groups, the response of the local LFPR, absent any demographic adjustments, may include both the direct cyclical effect as well as the effect of the migration response. However, national LFPR cyclicity would only contain the first effect, assuming that shocks do not induce sizeable migration out of the country. The response of the age-adjusted LFPR, though, would be comparable to national LFPR cyclicity, since it would not be affected by the migration channel.

To estimate the effect of a business-cycle shock on the composition of the state's population, we estimate [Equation 1](#) with the outcome $y_{s,t+k}$ being the log population of a single-year-age group in state s in period $t + k$.¹⁹ 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.²⁰

¹⁹Relative to [Amior and Manning \(2018\)](#), our analysis focuses on single-year-age groups instead of coarse age groups, estimates annual dynamic responses instead of decadal responses, and estimate the response to output shocks.

²⁰We use state-level data for the covered-area population for single-year-age groups from the U.S. Census Bureau. 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 domes-

Figure 6: 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 $t - 1$. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Regressions are weighted by population.
Source: BLS, BEA, and authors' calculations.

A negative business cycle shock causes the population between the ages of 25 and 40 to persistently decline in states exposed to the shock relative to those states without a shock (see Figure 6, 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 percent.

Upon impact of the shock, the migration response is still small, with essentially

tic) 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 covered-area population includes active members of the armed forces as well as 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.

all point estimates at 0 percent, but as time goes by, 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 changes in a state's population caused by a negative business cycle shock are entirely driven by people 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 net-population response might be concentrated in individuals ages 25 to 39, although formally testing these theories is outside the scope of our paper. First, people 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, people from other 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 6, as initially, the largest response is for people in their mid-20s. However, as time goes by and people get older, the response shifts to the right of the age distribution.²¹

²¹Our results also show a small *increase* in the population 17 to 22 year olds. Although testing the reasons behind the increase for this college-age group is beyond the scope of this paper, one plausible mechanism is that recessions cause reductions in income and wealth that make young people more likely to stay in state for their college education with more affordable tuition.

7 Differences in Long-Lived Cyclicity Across Groups

Business cycles can have different effects on different demographic groups. In this section, we examine how the cyclicity of the LFPR varies across the age, gender, education, and race/ethnicity distributions. Comparing young workers to older workers, men to women, and less-educated people to more-educated people, we find the LFPR for each former group is both more cyclical and features longer-lived cyclicity. These differences in long-lived cyclicity may create differential benefits for these groups from “running the economy hot” in years 5 to 7 after a shock, when the unemployment rate has fully recovered but the LFPR is still recovering (Aaronson et al., 2019).

7.1 Age

The labor market performance of prime-age people (ages 25 to 54) is often used as a benchmark for the cyclical state of the labor market as a whole. Understanding the cyclical response for the prime-age group is of considerable interest, as prime-age people make up 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 (2020) and Coglianesi (2018)), but there has been relatively less work on identifying the cyclical response of those variables from their long-run declining trends.²²

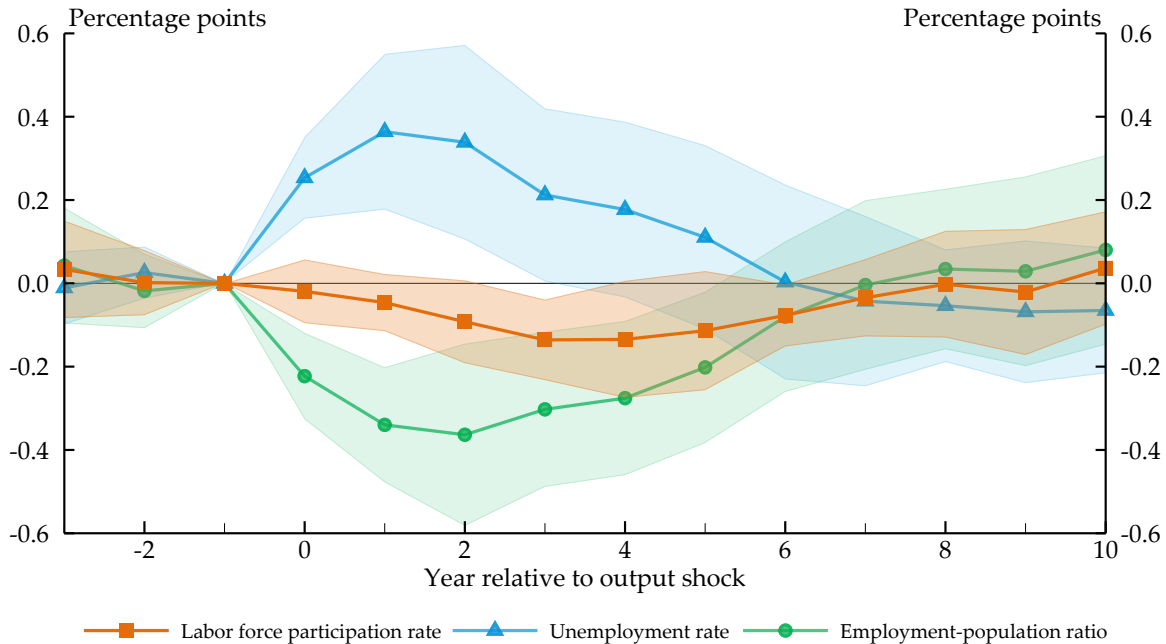
The cyclical response of the prime-age LFPR is similar to the overall response, albeit a bit smaller in magnitude. Figure 7 shows the estimated impulse response for the prime-age LFPR, along with unemployment rate and EPOP.²³ The LFPR declines steadily after the shock until it reaches its trough four years after the initial shock—well after the unemployment rate peaks—at about 0.14 percentage point below its pre-shock level, before gradually recovering and reaching its pre-shock level in year eight.

Compared to prime-age people, the LFPR for younger people (ages 16 to 24)

²²Although the main purpose of Aaronson et al. (2014b) and Montes (2018) is to build a forecasting model of the overall LFPR, both papers provide some evidence on the cyclicity 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.

²³Unlike our baseline results, we do not use age-adjusted participation rates for these subgroups. However, the results are very similar if we age-adjust the LFPRs within each age range. This is a consequence of the fact that changes in the demographic composition of the population mainly reflect changes across these age groups, rather than changes within them.

Figure 7: Cyclicity for Ages 25 to 54



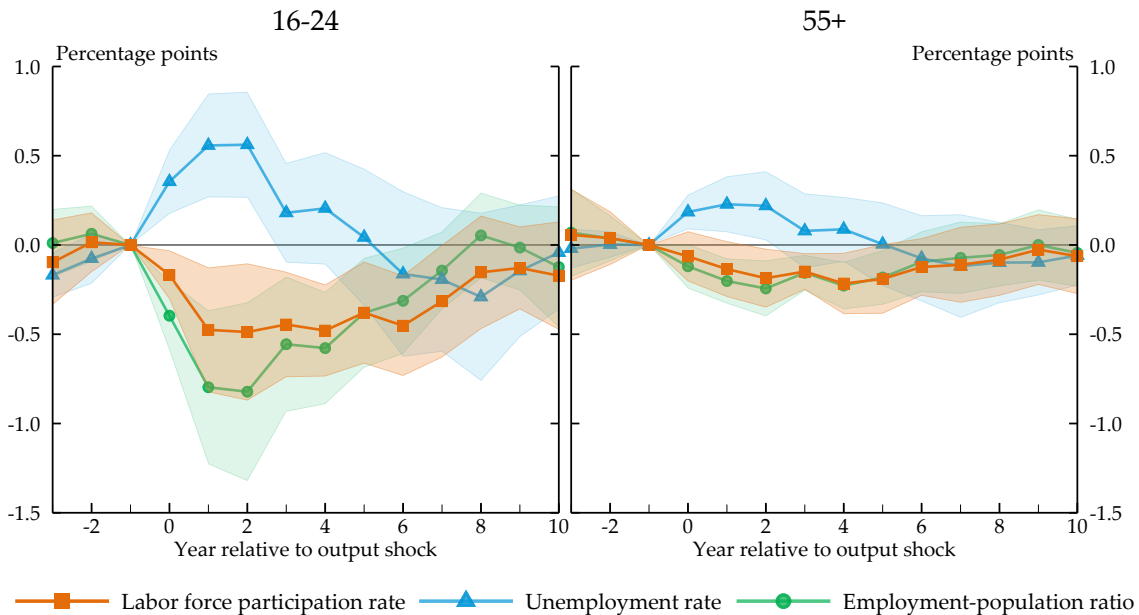
Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. F-statistic: 151.6. Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, own calculations.

responds more quickly and with a larger amplitude—reaching a trough of about -0.5 percentage point—but remains near its trough for many years and begins recovering later (Figure 8, left panel). The point estimate of the LFPR of younger people never fully recovers, as it settles at about 0.2 percentage point below its pre-shock value, although the upper end of the confidence interval suggests we cannot rule out a full recovery. The delayed recovery of the LFPR for younger people likely reflects the increase in time spent in schooling documented in Section 5.

The LFPR response for older people is similar to the response of the overall population, reaching its trough at about 0.2 percentage point four years after the shock (Figure 8, right panel). The LFPR for older people then begins to steadily recover 5 years after the shock and does not fully recovery until 9 years after the shock. For this age group, the shortfall of participation at its trough is likely due to higher rates of illness and disability, with no increase in retirements.

Figure 8: Cyclicity for Ages 16-24 and 55+



Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. F-statistic: 179.1 for 16–24, 162.6 for 55+. Regressions control for state and year fixed effects and are weighted by population.

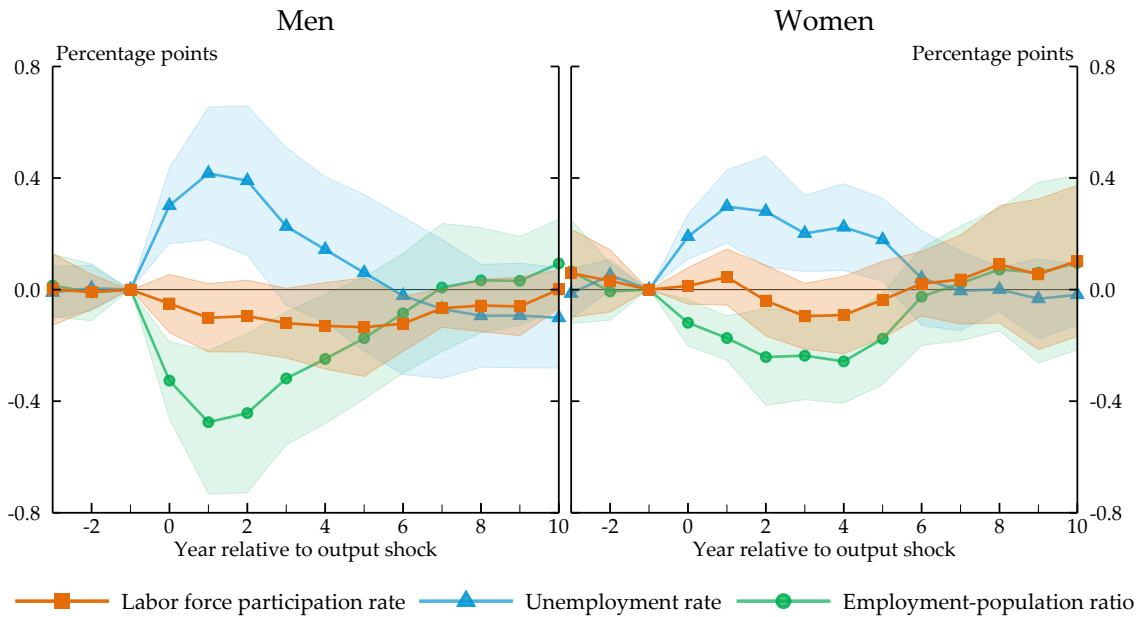
Source: BLS, BEA, own calculations.

7.2 Gender

Digging deeper into the prime-age LFPR responses, our results suggest that while both men and women have strong cyclicity, the magnitudes and timing of their responses are quite different. For men, the initial point estimate response shown in the left panel of Figure 9 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 0.15 percentage point at its trough 6 years after the shock. Although the confidence bands around those estimates are large due to the smaller sample sizes from splitting the prime-age group by gender, the decline in 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 point below its initial value. This rate fully recovers by about 6 years after the shock and settles at a rate slightly above its pre-shock value. Of course, the confidence bands around the

Figure 9: Cyclicity for Ages 25 to 54 by Sex



Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GDP growth in year 0. F-statistic: 151.1 for men, 152.1 for women. Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, own calculations.

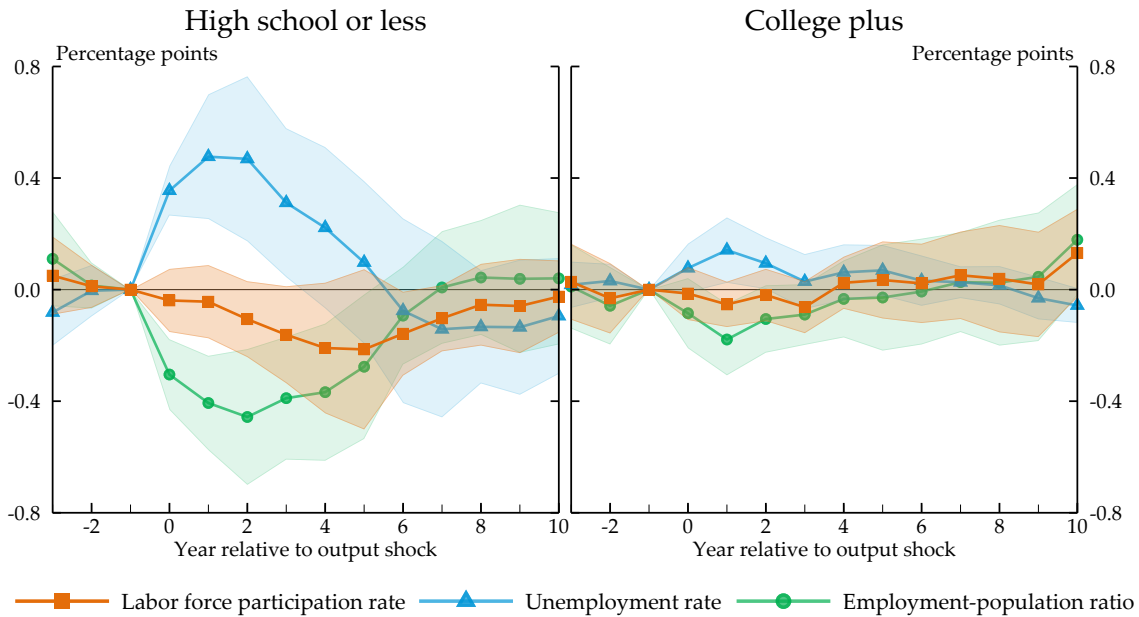
estimates for prime-age women are quite large, possibly due to large non-cyclical variation in the LFPR for prime-age women, and so one cannot reject the possibility that the LFPR of prime-age women does not respond to the shock at all.

7.3 Education

Labor market outcomes over at least the past 40 years have been quite different for less- and more-educated people. 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 less-educated people have been declining steadily over the past several decades, while the LFPR and EPOP for more-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, per-

Figure 10: Cyclicality for Ages 25 to 54 by Education



Note: Each line shows the estimated coefficients from Equation 1 for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GDP growth in year 0. F-statistic: 167.9 for high school degree or less, 137.7 for college degree or more. Individuals with some college but less than a four year degree are omitted. Regressions control for state and year fixed effects and are weighted by population. *Source:* BLS, BEA, own calculations.

haps, by changes in technology and globalization. Thus, to isolate cyclicality one needs to control for these long-term structural declines. Our approach using state-level business cycles and controlling for these national and international trends is well suited to isolate the effects of the business cycle and explore how they differ across education groups.

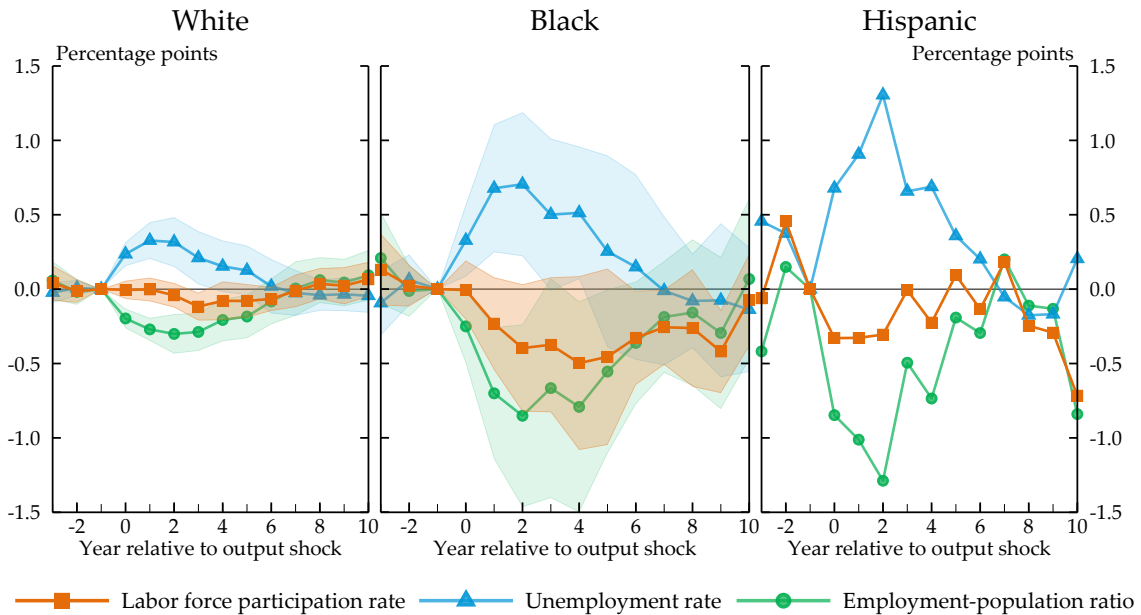
We find a starkly different evolution of the LFPR after a shock for less-educated prime-age workers compared to those with at least college degrees. For workers with a high school degree or less, the shock leads to a slow decline of the LFPR for about 5 years, reaching a trough of about 0.25 percentage point, before recovering subsequently. In contrast, workers with a college degree experience essentially no variation in LFPR following a shock.²⁴ This disparity is also found in the responses of the unemployment rate and EPOP, each of which respond substantially among the less-educated group but barely at all among the more-educated group.

²⁴We 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.

7.4 Race and Ethnicity

We also investigate the inequality of long-lived LFPR cyclicalities across race and ethnicity. As has been noted by [Cajner et al. \(2017\)](#) and others, business cycles are more costly for minority groups. We divide prime-age people in the CPS into racial and ethnic groups and estimate [Equation 1](#) for each group, showing the results in [Figure 11](#).

Figure 11: Cyclicalities for Ages 25 to 54 by Race/Ethnicity



Note: Each line shows the estimated coefficients from [Equation 1](#) for the associated labor market outcome. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Individuals not reporting either white, Black, or Hispanic are omitted. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GDP growth in year 0. F-statistic: 186.5 for white, 134.9 for Black, 10.4 for Hispanic. Confidence interval for Hispanic not shown due to low F-statistic. Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, own calculations.

We find that shocks lead to larger and more long-lived declines in LFPR among minority groups. While the white LFPR falls by only 0.1 percentage point after a shock, the Black LFPR falls by 0.5 percentage point. The Black LFPR remains depressed for substantially longer, and only fully recovers ten years after the shock, well after the white LFPR has recovered. The responses for Hispanic workers are also large, although our results for this group are much noisier due to a lower-powered instrument when weighting states by the Hispanic population.

8 What Drives the Shocks?

We examine what drives the variation in our output shock. We find similar responses to contractionary and expansionary shocks, suggesting that our effects are not being driven by asymmetries. More of our variation comes from the pre-1994 period, with estimates using only post-1994 data being similar overall but substantially noisier. The variation in the shift-share instrument is driven by a handful of industries including motor vehicle production, oil and gas extraction, securities and commodities brokers, and farms, but our estimated effects are similar if these industries are excluded. Further, we show that our shocks primarily reflect short-lived shocks to productivity growth, which then spill over to persistent effects on employment. Overall, we find that our results are not being driven by a single source of variation, and instead reflect common responses to shocks in a wide variety of environments.

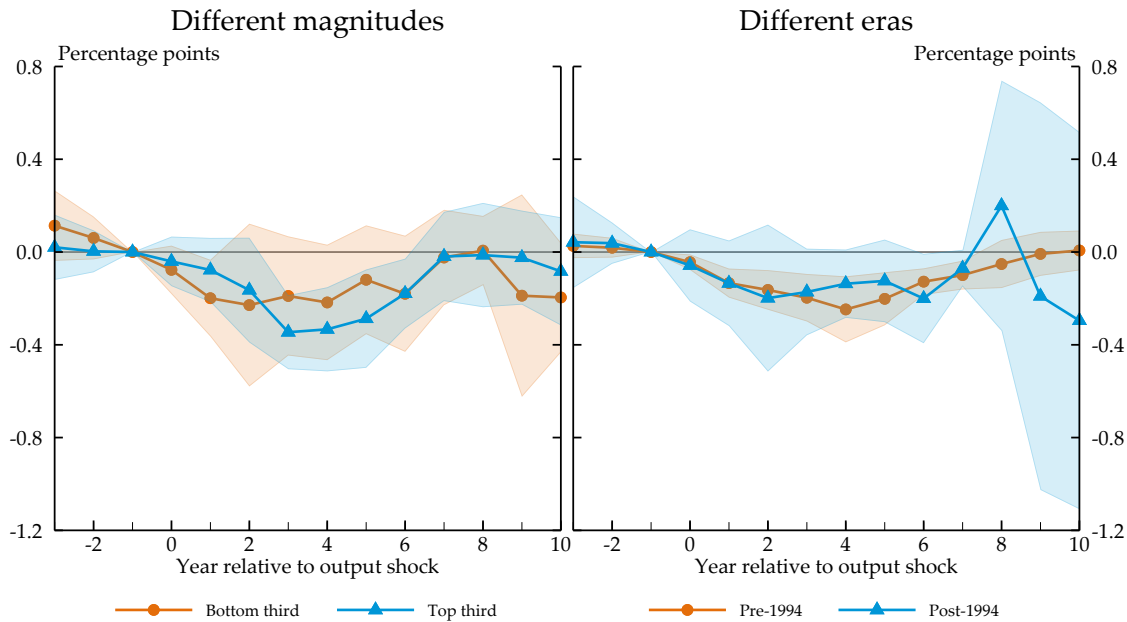
8.1 Expansions vs. Contractions

Our estimated impulse responses are an average of the effects of expansionary and contractionary shocks, which may not be informative if these effects are starkly different. To examine whether expansionary and contractionary shocks have different effects, we divide the distribution of shocks into thirds and estimate the impulse responses separately for each third. In the left panel of [Figure 12](#), we present the effects of expansionary shocks (top third) and contractionary shocks (bottom third), normalizing both to show the effect of a negative 1 percentage point shock. Both impulse responses have similar patterns, and we cannot reject that the two are the same. This result suggests that our baseline estimates, which combine the response of both expansionary and contractionary shocks, are a reasonable guide for a wide range of shocks.

8.2 Differences over Time

Our instrument also combines variation over time, including periods with different macroeconomic dynamics. Business cycles since 1990 have been characterized by jobless recoveries ([Jaimovich and Siu, 2020](#)), while earlier periods included more rapid recoveries in the labor market. Additionally, our CPS sample includes data both before and after the 1994 redesign, which substantially changed how the survey was collected.

Figure 12: Cyclical Responses to Different Types of Shocks



Note: Each line shows the estimated coefficients from Equation 1 for the LFPR, using only the specified sample of shocks. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. In all specifications, the LFPR is adjusted for changes in the age-by-sex composition of the population. F-statistics: 51.7 (bottom-third), 45.7 (top-third), 146.4 (pre-1994), 31.2 (post-1994). Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, and authors' calculations.

To test whether the cyclical nature of the LFPR has changed over time, we divide our sample into pre- and post-1994 periods. For each period, we separately estimate the impulse response and plot these estimates in the right panel of Figure 12. Although the post-1994 estimates are substantially noisier, the two point estimates are similar and we cannot rule out that the two are the same. This suggests that most of the variation in the instrument in our baseline estimates comes from the earlier period, but it does not exclusively drive our estimates.

8.3 Decomposing the Shift-Share Instrument

To further examine where the variation in our shift-share instrument comes from, we decompose the variation using the approach of Goldsmith-Pinkham, Sorkin and Swift (2020). For simplicity, we focus on the response of the LFPR four years after the shock, which is the point that it reaches its trough in our main estimates. To

Table 1: Rotemberg weights in GSP Shift-Share Instrument

(a) By Industry/Year				(b) By Industry			(c) By Year		
Industry	Year	α_{kt}	β_{kt}	Industry	α_k	β_k	Year	α_t	β_t
Oil & gas	1986	0.13	0.40	Motor vehicles	0.30	0.09	1980	0.22	0.16
Oil & gas	1980	0.12	0.24	Oil & gas	0.28	0.38	1986	0.17	0.35
Securities	2009	0.08	0.10	Securities	0.12	0.11	1983	0.10	0.20
Motor vehicles	2010	0.07	0.07	Farms	0.06	0.13	2009	0.10	-0.02
Motor vehicles	1980	0.05	0.09	Primary metals	0.03	-0.19	2010	0.07	0.12
Oil & gas	1981	0.04	0.35	Computers & electronics	0.02	0.04	1982	0.07	0.14
Motor vehicles	2009	0.03	-0.05	Trans. eq. excl. motor veh.	0.02	0.20	1992	0.04	0.18
Motor vehicles	1983	0.03	0.07	Federal govt. - military	0.02	0.60	2001	0.04	0.50
Oil & gas	1983	0.03	0.21	State & local govt.	0.02	0.29	1994	0.03	0.04
Motor vehicles	1992	0.02	0.45	Chemicals	0.01	0.11	1981	0.03	0.21
All other	All other	0.38	0.16	All other	0.12	0.18	All other	0.14	0.22

Note: Tables show the Rotemberg weights for the GSP shift-share instrument used in our main estimates. Each panel shows the top 10 Rotemberg weights in each category, along with the total among all non-top-10 entries. Outcome is the change in the LFPR four years after the shock; the total effect is equal to 0.19 in our main specification using the non-leave-one-out version of the instrument. *Source:* BLS, BEA, and authors' calculations.

compute the Rotemberg weights for each industry-year pair, we compute

$$\hat{\alpha}_{kt} = \frac{g_{kt} \zeta'_{kt} \Delta \text{GSP}_{t,t-1}^{\perp}}{\sum_{k'} \sum_{t'} g_{k't'} \zeta'_{k't'} \Delta \text{GSP}_{t',t'-1}^{\perp}}, \quad \hat{\beta}_{kt} = \frac{\zeta'_{kt} \Delta \text{LFPR}_{t+4,t-1}^{\perp}}{\zeta'_{kt} \Delta \text{GSP}_{t,t-1}^{\perp}}, \quad \hat{\beta} = \sum_k \sum_t \hat{\alpha}_{kt} \hat{\beta}_{kt} \quad (10)$$

where $\Delta \text{GSP}_{t,t-1}^{\perp}$ is GSP growth and $\Delta \text{LFPR}_{t+4,t-1}^{\perp}$ is the cumulative change in LFPR by four years after shock, both residualized on state and year fixed effects, ζ'_{kt} is the lagged industry share for industry k in year t , and g_{kt} is the national growth rate of industry k in year t . We depart from our baseline specification in using the national growth rate for g_{kt} , instead of the leave-one-out growth rate, in order to align with the calculation of Rotemberg weights.²⁵

Importantly, we treat each industry and year as a distinct instrument, using the variation from the shares to identify each effect. Our baseline estimate is a weighted average of these effects, where the weights are the Rotemberg weights outlined above. An alternative interpretation of the shift-share instrument is that the variation comes from the industry shocks, as outlined in [Borusyak, Hull and Jaravel \(forthcoming\)](#).

Much of the variation in the shift-share instrument comes from a small number of industry-year instruments. Panel (a) of [Table 1](#) shows the top 10 industry-year instruments, along with their weights $\hat{\alpha}_{kt}$ and estimated effects $\hat{\beta}_{kt}$. The instruments

²⁵Our baseline results are little changed using the national growth rate instead of the leave-one-out growth rate.

contributing the most weight include shocks to oil & gas extraction during the 1980s, as well as shocks to motor vehicle production and securities during recessions. Collectively, the top 10 instruments account for about 62 percent of the total weight. Most of the shocks have estimated β s close to our main estimate, including the total of shocks outside the top 10. In this way, no single shock drives our result.

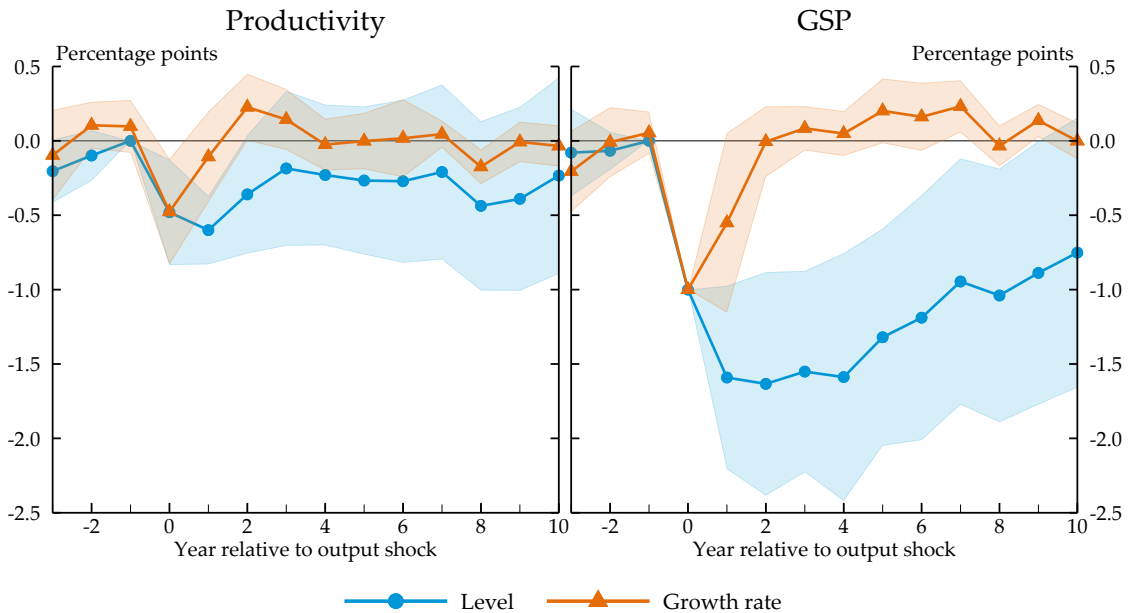
We also aggregate the weights to show the most important industries, pooling across time periods, and the most important time periods, pooling across industries. Panel (b) of [Table 1](#) shows that 3/4 of the shift-share instrument variation comes from just four industries—motor vehicle production, oil & gas extraction, securities & commodities brokers, and farms. Nonetheless, these industries do not exclusively drive our result, as the estimated effect pooling across all other industries is 0.18, very close to our baseline estimate. Panel (c) of [Table 1](#) shows that our instrument derives a substantial amount of variation from recessions, with the top 10 years including at least one year from each of the five national recessions that took place during our sample period, but also includes variation from non-recessionary years. Almost all years have coefficients close to our baseline estimate, indicating that our estimates are not being driven by a single year or recession.

8.4 Effects on Productivity

Our shocks to output could result either from lower output per worker, or fewer workers, or some combination thereof. We have shown in our baseline estimates that employment declines, but some of the output effect could still be driven by labor productivity—defined here as GSP per employee. Importantly, the potential for our instrument to contain variation in productivity shocks sets it apart from shift-share instruments that are based purely on employment.

[Figure 13](#) shows the estimated impulse response of productivity to a temporary negative 1 percentage point output shock, using the same approach as in [Equation 1](#). The left panel shows the effect on yearly growth rates of productivity, along with the cumulated effect on the level of productivity. Productivity grows by about 0.5 percentage point less in the year when the shock takes place, but grows similarly afterwards. This leads to a level of productivity that is permanently about 0.25–0.5 percent lower after the shock than before. Productivity accounts for about half of the initial shock to output (shown in the right panel of [Figure 13](#)), with the remainder accounted for by employment. As productivity is stable after the initial shock, the further decline in output in year 1 and the subsequent partial recovery entirely reflect employment. This points to output shocks being initially driven by productiv-

Figure 13: Effects of Shocks on Productivity and Output



Note: Each line shows the estimated coefficients from Equation 1 for the specified outcome, either in levels relative to year -1 or in growth rates. The bands around each line show a 95% confidence interval, based on standard errors clustered by state. The left panel shows the response of real productivity, defined as real GSP per worker. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population.
Source: BLS, BEA, and authors' calculations.

ity before employment adjusts in response, with time aggregation leading to some of this response appearing in the same year as the shock. These estimates also indicate that our instrument picks up an important source of variation—productivity shocks—which would be omitted in an employment-based shift-share instrument.

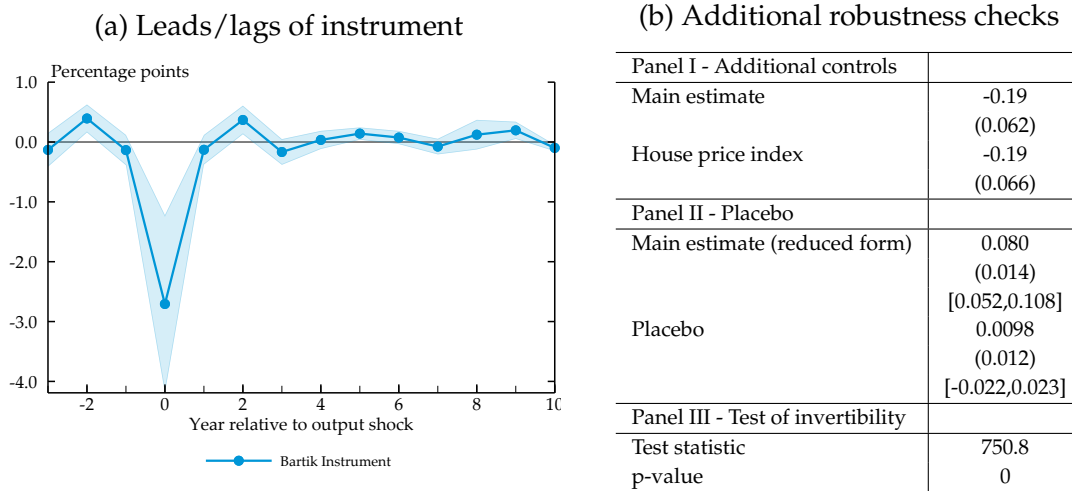
9 Robustness

In this section, we show several robustness checks for our methodology.

9.1 Lead-Lag Exogeneity

One of the conditions required for our research design to identify the impulse response of the LFPR is that the instrument satisfies lead-lag exogeneity, as laid out in Equation 6 (Stock and Watson, 2018). A necessary, though not sufficient, condition for lead-lag exogeneity is that the instrument should be uncorrelated with leads and

Figure 14: Robustness Checks



Note: In the left panel, the line shows the estimated coefficients from Equation 1 using the shift-share instrument as the outcome, and the band around the line shows a 95% confidence interval, based on standard errors clustered by state. The right panel shows the estimated response of the age-adjusted LFPR four years after a shock (panels I and II), as well as the results of the Stock and Watson (2018) test of invertibility. Standard errors clustered by state are shown in parentheses. In panel I, coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. In panel II, the 95% confidence interval is shown in brackets; for the placebo specification this is the empirical confidence interval taken from the 2.5th percentile to the 97.5th percentile across placebo estimates. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population.

Source: BLS, BEA, FHFA, and authors' calculations.

lags of itself, which we can test empirically. Given that our instrument is based on industry growth rates and shares, which can be persistent over time, there is some potential for the instrument to be correlated with leads and lags of itself.

To examine whether our instrument is correlated with its leads and lags, we estimate Equation 1 using our shift-share instrument as the outcome variable. This impulse response is reported in the left panel of Figure 14. The coefficient in period 0, 2.71, is the inverse of our first stage coefficient, γ , and is highly statistically significant as a result. Importantly, though, all of the other coefficients are close to zero and almost all of them are statistically indistinguishable from zero.

9.2 Controlling for House Price Growth

To verify the robustness of our results, we show that they remain unaffected when controlling for local house price growth. Our focus in this paper is on the response of LFPR to changes in output, which we take to represent changes in the production process. An alternative reason that measured output can change is if capital income

changes unrelated to current production (for example, through home price appreciation).²⁶ Controlling for home price appreciation as measured by the state-level FHFA price index in the period before the shock addresses this concern. Panel I of [Figure 14](#) shows that our estimates are little changed from the baseline if we control for home price growth.

9.3 Placebo

We cluster our standard errors at the state level in our baseline estimates, but [Adão, Kolesár and Morales \(2019\)](#) point out that this may be insufficient in some circumstances. Our instrument exploits variation across places with different industry exposure, and the residuals for states with similar industry exposure may be correlated. Our clustering approach does not exactly capture this structure, raising a concern that our standard errors may be incorrect.

We examine the relevance of this critique for our setting using a placebo exercise similar to the one proposed by [Adão, Kolesár and Morales \(2019\)](#). In place of our shift-share instrument, we estimate the reduced-form version of our main specification using a placebo shift-share instrument, where the national growth rates of each industry have been replaced with random draws from a normal distribution with the same mean and variance as the observed growth rates. We repeat this procedure 100 times, obtaining a placebo estimate for each, and report the distribution of these placebo estimates along with our baseline in panel II of [Figure 14](#). Unlike the cases examined by [Adão, Kolesár and Morales \(2019\)](#), we find that the spread of placebo estimates is similar to or a bit smaller than the confidence intervals obtained from standard errors clustered at the state level. This result suggests that our approach to inference is valid, and if anything is a bit conservative.

9.4 Local Projections vs. VAR

A key departure of our approach from the literature is the use of local projections regressions instead of a VAR to estimate impulse response functions. Both [Blanchard and Katz \(1992\)](#) and [Dao, Furceri and Loungani \(2017\)](#) use VAR methods to estimate impulse responses and find roughly similar cyclical timing for the unemployment rate and LFPR. However, VAR methods can fail to identify the correct impulse responses even when the instrument conditions are met if the impulse responses are

²⁶In the opposite direction, [Mian and Sufi \(2014\)](#) show that house price declines during the Great Recession led to lower employment through deteriorating household balance sheets.

not invertible, but local projections do not require this assumption for identification (Stock and Watson, 2018).

To test whether VAR methods are appropriate for our setting, we conduct a test of invertibility following Stock and Watson (2018). This is a Hausman (1978)-type test, where, under the null hypothesis of invertibility, both methods should deliver similar estimates but with VAR estimates more efficient, while under the alternative they would return different estimates. We report the test statistic in panel III of Figure 14 along with the associated p-value. We are able to strongly reject the null hypothesis of invertibility, implying that local projections are the only suitable method for examining the cyclical nature of LFPR with our approach.

10 Conclusion

We estimate the effect of a business cycle shock on the LFPR and show that the LFPR is cyclical, but it responds with a smaller elasticity, a more delayed impact, and a longer recovery than the unemployment rate. Our approach uses state-level variation in business cycles to estimate the cyclical nature 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 cyclical nature 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.

Our results indicate that measuring labor market slack requires looking beyond the unemployment rate. While traditional views hold that the unemployment rate is a sufficient statistic for slack, the long-lived cyclical nature of the LFPR poses problems for this view. During the period 5 to 7 years after a shock, the unemployment rate has essentially fully recovered, but the LFPR still has room to rise before it returns to its pre-shock level. Observers who focus solely on the unemployment rate during this period will incorrectly conclude that the economy has reached full employment, when in fact employment is still below potential.

A complete view of labor market slack requires examining the LFPR in addition to the unemployment rate, and perhaps may go further to include the differential cyclical nature of different demographic groups. Long-lived cyclical nature is especially

prone among younger workers, men, less educated workers, and racial and ethnic minorities, each of which is also more exposed to business cycles in the form of unemployment. Our results indicate that these groups have the most to gain from maintaining business cycle recoveries until the LFPR has fully recovered, and also the most to lose if long-lived cyclicalities in the LFPR is ignored.

References

- Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan.** 2014a. "Declining Labor Force Participation and its Implications for Unemployment and Employment Growth." *Economic Perspectives*, 38(4).
- Aaronson, Stephanie R., Mary C. Daly, William L. Wascher, and David W. Wilcox.** 2019. "Okun Revisited: Who Benefits Most from a Strong Economy?" *Brookings Papers on Economic Activity*, 2019(1): 333–404.
- Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher.** 2014b. "Labor Force Participation: Recent Developments and Future Prospects." *Brookings Papers on Economic Activity*, 2014(2): 197–275.
- Abowd, John M., and Arnold Zellner.** 1985. "Estimating Gross Labor-Force Flows." *Journal of Business & Economic Statistics*, 3(3): 254–283.
- Abraham, Katharine G., and Melissa S. Kearney.** 2020. "Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence." *Journal of Economic Literature*, 58(3): 585–643.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales.** 2019. "Shift-Share Designs: Theory and Inference." *The Quarterly Journal of Economics*, 134(4): 1949–2010.
- Amior, Michael, and Alan Manning.** 2018. "The Persistence of Local Joblessness." *American Economic Review*, 108(7): 1942–70.
- Balakrishnan, Ravi, Mai Dao, Juan Sole, and Jeremy Zook.** 2015. "Recent U.S. Labor Force Dynamics; Reversible or not?" International Monetary Fund IMF Working Papers 15/76.
- Ball, Laurence, Daniel Leigh, and Prakash Loungani.** 2017. "Okun's Law: Fit at 50?" *Journal of Money, Credit and Banking*, 49(7): 1413–1441.
- Bartik, Timothy J.** 1991. *Who Benefits from State and Local Economic Development Policies?* Books from Upjohn Press, W.E. Upjohn Institute for Employment Research.
- Beraja, Martin, Erik Hurst, and Juan Ospina.** 2019. "The Aggregate Implications of Regional Business Cycles." *Econometrica*, 87(6): 1789–1833.
- Blanchard, Olivier Jean, and Lawrence F. Katz.** 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 1992(1): 1–76.

- Borusyak, Kirill, Peter Hull, and Xavier Jaravel.** forthcoming. "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies*.
- Bound, John, and Harry J. Holzer.** 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics*, 18(1): 20–54.
- Bradbury, Katharine L.** 2000. "Rising Tide in the Labor Market: To What Degree Do Expansions Benefit the Disadvantaged?" *New England Economic Review*, 3–34.
- Cadena, Brian C., and Brian K. Kovak.** 2016. "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession." *American Economic Journal: Applied Economics*, 8(1): 257–90.
- Cairó, Isabel, Shigeru Fujita, and Camilo Morales-Jiménez.** 2021. "The Cyclicalities of Labor Force Participation Flows: The Role of Labor Supply Elasticities and Wage Rigidity." *Review of Economic Dynamics*.
- Cajner, Tomaz, Tyler Radler, David Ratner, and Ivan Vidangos.** 2017. "Racial Gaps in Labor Market Outcomes in the Last Four Decades and over the Business Cycle." FEDS Working Paper.
- CBO.** 2007. "The Budget and Economic Outlook: 2021 to 2031." Congressional Budget Office.
- Chodorow-Reich, Gabriel.** 2019. "Geographic cross-sectional fiscal spending multipliers: What have we learned?" *American Economic Journal: Economic Policy*, 11(2): 1–34.
- Coglianesi, John.** 2018. "The Rise of In-and-Outs: Declining Labor Force Participation of Prime Age Men." Working Paper.
- Council of Economic Advisers.** 2014. "The Labor Force Participation Rate since 2007: Causes and Policy Implications." mimeo.
- Dao, Mai, Davide Furceri, and Prakash Loungani.** 2017. "Regional Labor Market Adjustment in the United States: Trend and Cycle." *The Review of Economics and Statistics*, 99(2): 243–257.
- Davidson, Russell, and James MacKinnon.** 1993. "Estimation and Inference in Econometrics." Oxford University Press.

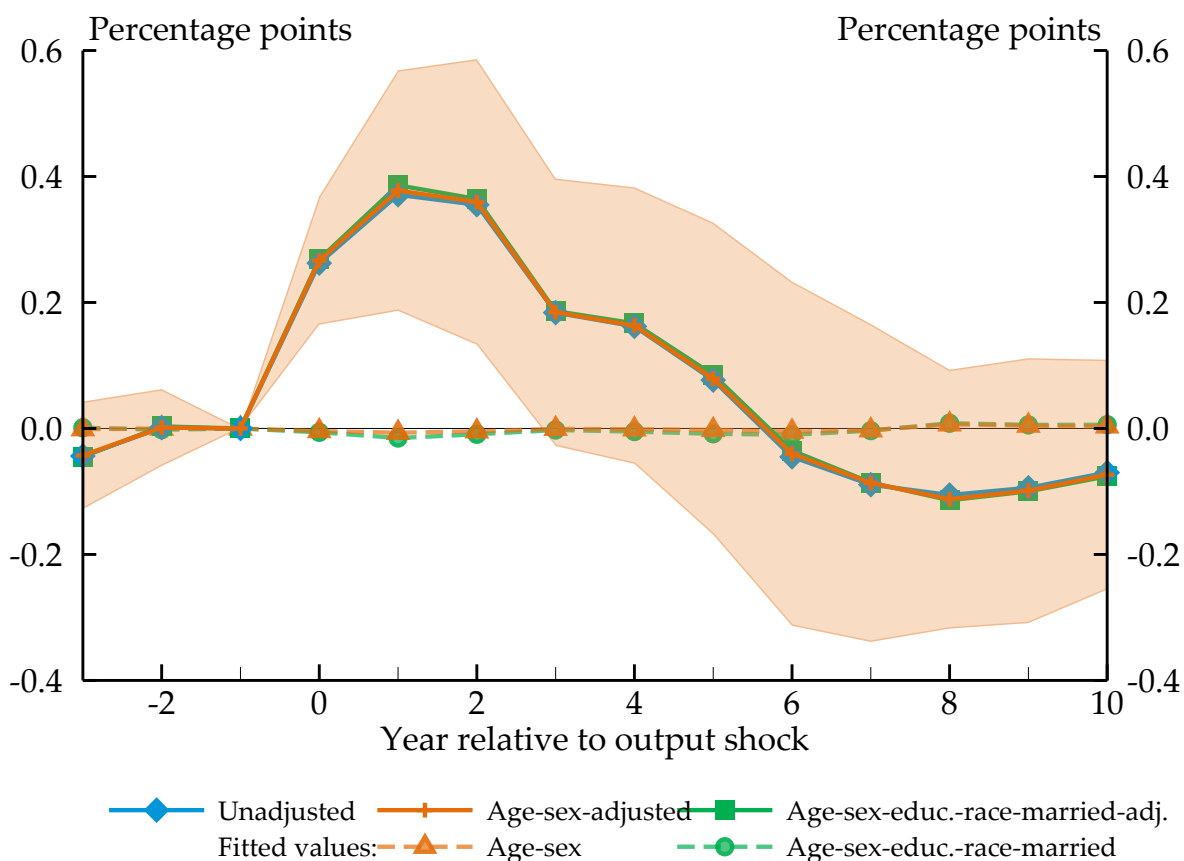
- Decressin, Jörg, and Antonio Fatas.** 1995. "Regional Labor Market Dynamics in Europe." *European Economic Review*, 39(9): 1627–1655.
- Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson.** 2019. "A Plucking Model of Business Cycles." NBER Working Paper 26351.
- Elsby, Michael, Bart Hobijn, Fatih Karahan, Gizem Koşar, and Ayşegül Şahin.** 2019. "Flow Origins of Labor Force Participation Fluctuations." *AEA Papers and Proceedings*, 109: 461–64.
- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin.** 2015. "On the Importance of the Participation Margin for Labor Market Fluctuations." *Journal of Monetary Economics*, 72: 64–82.
- Erceg, Christopher J., and Andrew 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.
- Fallick, Bruce, and Pawel Krolkowski.** 2019. "Excess Persistence in Employment of Disadvantaged Workers." Federal Reserve Bank of Cleveland Working Papers 201801R.
- Fukui, Masao, Emi Nakamura, and Jón Steinsson.** 2018. "Women, wealth effects, and slow recoveries." National Bureau of Economic Research.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review*, 110(8): 2586–2624.
- Hausman, J. A.** 1978. "Specification Tests in Econometrics." *Econometrica*, 46(6): 1251–1271.
- Herbst, Edward, and Benjamin K. Johansson.** 2020. "Bias in Local Projections." Board of Governors of the Federal Reserve System (U.S.) Finance and Economics Discussion Series 2020-010.
- Hershbein, Brad, and Bryan A Stuart.** 2020. "Recessions and local labor market hysteresis."
- Hornbeck, Richard, and Enrico Moretti.** 2018. "Who Benefits from Productivity Growth? Direct and Indirect Effects of Local TFP Growth on Wages, Rents, and Inequality." NBER Working Paper 24661.

- Hornstein, Andreas, and Marianna Kudlyak.** 2019. "Aggregate Labor Force Participation and Unemployment and Demographic Trends." Federal Reserve Bank of San Francisco Working Paper Series 2019-7.
- Hotchkiss, Julie L., and Robert Elijah Moore.** 2018. "Some Like it Hot: Assessing Longer-Term Labor Market Benefits from a High-Pressure Economy." Andrew Young School of Policy Studies Research Paper Series 18-01.
- Howard, Greg.** 2020. "The Migration Accelerator: Labor Mobility, Housing, and Demand." *American Economic Journal: Macroeconomics*, 12(4): 147–79.
- Hoynes, Hilary, Douglas L Miller, and Jessamyn Schaller.** 2012. "Who suffers during recessions?" *Journal of Economic Perspectives*, 26(3): 27–48.
- Hoynes, Hilary Williamson.** 2000. "Local Labor Markets and Welfare Spells: Do Demand Conditions Matter?" *Review of Economics and Statistics*, 82(3): 351–368.
- Jaimovich, Nir, and Henry E. Siu.** 2020. "Job Polarization and Jobless Recoveries." *The Review of Economics and Statistics*, 102(1): 129–147.
- Jefferson, Philip N.** 2008. "Educational Attainment and the Cyclical Sensitivity of Employment." *Journal of Business & Economic Statistics*, 26(4): 526–535.
- Jordà, Òscar.** 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.
- Juhn, Chinhui, and Simon Potter.** 2006. "Changes in labor force participation in the United States." *Journal of Economic Perspectives*, 20(3): 27–46.
- Krueger, Alan B.** 2017. "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity*, 2017(2): 1–87.
- Mian, Atif, and Amir Sufi.** 2014. "What explains the 2007–2009 drop in employment?" *Econometrica*, 82(6): 2197–2223.
- Monras, Joan.** 2018. "Economic Shocks and Internal Migration." CEPR Discussion Paper No. DP12977.
- Montes, Joshua.** 2018. "CBO's Projection of Labor Force Participation Rates." Congressional Budget Office Working Papers 2018-04.

- Nakamura, Emi, and Jon Steinsson.** 2014. "Fiscal stimulus in a monetary union: Evidence from US regions." *American Economic Review*, 104(3): 753–92.
- Nakamura, Emi, and Jón Steinsson.** 2018. "Identification in macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.
- Okun, Arthur M.** 1973. "Upward Mobility in a High-Pressure Economy." *Brookings Papers on Economic Activity*, 1973(1): 207–262.
- Olea, José Luis Montiel, and Mikkel Plagborg-Møller.** 2020. "Local Projection Inference is Simpler and More Robust Than You Think." Working Paper.
- Plagborg-Møller, Mikkel, and Christian K. Wolf.** 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.
- Shimer, Robert.** 1999. "Why is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual 1998, volume 13*. 11–74. MIT Press.
- Staiger, Douglas, and James H. Stock.** 1997. "Instrument Variables Regression with Weak Instruments." *Econometrica*, 65(3): 557–586.
- Stock, James H., and Mark W. Watson.** 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.
- Weinstein, Russell.** 2018. "Dynamic Responses to Labor Demand Shocks: Evidence from the Financial Industry in Delaware." *Journal of Urban Economics*, 106: 27–45.
- Wilson, Valerie.** 2015. "The Impact of Full Employment on African American Employment and Wages." Economic Policy Institute Report.
- Yagan, Danny.** 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy*, 127(5): 2505–2558.
- Yellen, Janet L.** 2014. "Labor Market Dynamics and Monetary Policy." Speech by Chair Janet L. Yellen at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming.

A Additional Results

Figure A.1: Unemployment Rate Cyclicity by Demographic Adjustment



Note: Each line shows the estimated coefficients from Equation 1 using the specified adjusted/unadjusted LFPR or fitted values as the outcome. The band around the orange solid line shows a 95% confidence interval, based on standard errors clustered by state. Coefficients are normalized to show the effect of a temporary -1 percentage point shock to GSP growth in year 0. F-statistic: 149.1. Regressions control for state and year fixed effects and are weighted by population. *Source:* BLS, BEA, own calculations.