

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

Tomaz Cajner John Coglianesi Joshua Montes

October 26, 2020

Preliminary, comments welcome.

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 from the response of the age-adjusted LFPR, and we show that it is problematic to use the unadjusted LFPR because local shocks spur changes in the population of high-LFPR age groups through the migration channel. LFPR cyclical 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 seminar participants at the Federal Reserve Board of Governors, WALES 2020, and SOLE 2020 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.

1 Introduction

How does the U.S. labor force participation rate respond to business cycle shocks? Many observers have noted that the labor force participation rate (LFPR)—the share of adults either working or looking for work—exhibits some degree of cyclicity (see, for example, [Aaronson et al., 2014b](#); [Council of Economic Advisers, 2014](#); [Montes, 2018](#)). Measuring the degree of cyclicity in the LFPR is complicated, though, by the presence of trend movements that reflect 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. Since many observers disagree about the exact magnitude of these trends, there is also substantial disagreement about the extent of cyclicity in labor force participation.

Our approach uses state-level variation in business cycles to estimate the cyclicity of labor force participation, sidestepping the issue of identifying trend changes in labor force participation at the national level. We measure cyclical variation at the state level using Gross State Product (GSP) and estimate the response of labor market outcomes using the local projections method ([Jordà, 2005](#)). Our research design identifies the response of labor market outcomes to unexpected declines in state 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 GSP 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 and takes about 8 years to complete a full recovery. Specifically, the LFPR declines slowly yet persistently following a negative 1 percent output shock and does not trough until 4 years after that shock—at about 0.2 percentage point below its initial value. This response is in contrast with that of the unemployment rate, which spikes quickly, peaks a year after a negative output shock, and has an elasticity that is about twice as large in absolute value compared to the LFPR.¹ By the time that the unemployment rate fully recovers 6 years after the shock, the LFPR is only in the early stages of its cyclical recovery. The LFPR does eventually complete its cyclical recovery, though, but not until about 8 years after the initial shock. This delay in recovery between the LFPR and the unemployment rate suggests that observers who focus only on the unemployment rate underestimate the

¹The unemployment rate peaks at 0.38 percentage point above its pre-shock level, which is just below the range of Okun’s law coefficients estimated in the literature ([Ball et al., 2017](#)), supporting that our method measures cyclicity accurately.

extent of slack remaining in the labor market after a recession, particularly in the period approximately 6 years or more after the initial shock.

Several possibilities could explain the long-lived cyclical dynamics where the LFPR takes several years to trough in response to a shock and nearly a decade to fully recover. For instance, it could represent hysteresis or discouraged worker effects in the labor market, as people who are laid off during a cyclical downturn become cut off from work, leave the labor force, and only return as the labor market tightens. These people may want a job but think that prospects for finding work are low, especially in the first few years following a shock. In this case, they may pursue out of the labor force activities that have persistent labor supply effects, such as increased schooling or taking care of the home, or methods of income replacement, such as early retirement or disability uptake.

We help distinguish between these possible explanations by estimating how peoples' reported reasons for being out of the labor force and their desire for a job respond to an exogenous shock to output. The share of people who are out of the labor force but want a job tracks the cyclical increase in total nonparticipation well for several years immediately following a shock, suggesting that worker discouragement over job prospects slowly but persistently builds after a shock. Further, the cyclical evolution of self-reported reasons for non-participation suggests that people leaving the labor force in response to a shock are engaging in productive activities, though activities with potentially persistent labor supply effects, as the decline in LFPR is largely driven by people who report being out of the labor force for home and family reasons or to attend school.

Our main approach controls for structural changes in the composition of state-level populations that are caused by local-level shocks and thus better isolates a true cyclical response to an output shock. Local shocks induce shifts in the composition of population through interstate migration that can cause permanent, structural changes to the levels of the LFPR across states, and these population effects have made it difficult to interpret the response of state-level LFPRs to output shocks in the context of the cyclicity of the national LFPR. To address this issue, we control for compositional changes in our main specification by estimating the response of an age-adjusted, state-level LFPR to an output shock. We adjust for age by residualizing the LFPR using person-level data to net out the composition component explained by the age distribution of the people in each state.

Nevertheless, we show that negative shocks do induce changes in the composition of the population that structurally lower the unadjusted LFPR in affected states. The permanently shortfall in the LFPR in those states is driven entirely by the compositional component of the LFPR, and in particular, by changes in the age composition of the population. More specifi-

cally, negative output shocks lead to a lower long-run level of the population for adults ages 25 to 39 years old—who tend to have higher participation rates than average—while the rest of the population is largely unaffected. This decline likely reflects out-migration by individuals in this age range. Our work adds to the understanding of the migratory adjustment mechanism of local shocks in the studies by [Blanchard and Katz \(1992\)](#) and [Dao et al. \(2017\)](#) work by showing that migration is driven predominantly by young, prime-age people.

We also document heterogeneity in the cyclical response of the LFPR across different demographic groups. The response of the prime-age LFPR shows the same delayed but full cyclical recovery of the LFPR—it reaches its trough 4 years after the shock and does not fully recover until 8 years out. We find that these patterns vary noticeably across gender, with the cyclical response of the LFPR much more long-lived for men than for women, though both groups eventually fully recover. Younger workers (ages 16–24) exhibit a much larger cyclical response of the LFPR than does the overall population, while older workers (ages 55+) show a lower degree of cyclicity. 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 cyclicity across racial and ethnic groups, with the prime-age black LFPR exhibiting larger, longer-lived cyclicity than the prime-age white LFPR.

Our paper is related to several strands of literature. First, several recent papers provide estimates of LFPR cyclicity and its structural trend ([Aaronson et al., 2014b,a](#); [Council of Economic Advisers, 2014](#); [Erceg and Levin, 2014](#); [Balakrishnan et al., 2015](#); [Montes, 2018](#); [Hornstein and Kudlyak, 2019](#)). While these papers do argue that the cyclical response of LFPR can be delayed, one of the main contributions of our paper is to use a method that is particularly well-suited for causally estimating long lags in LFPR cyclicity. More precisely, unlike the previous papers in this literature we estimate the dynamic response of LFPR to output shocks by using the local projection regressions, which allow for the possibility of very persistent effects on LFPR. Moreover, by using a shift-share instrumental variable approach, we are able to establish a causal link between business cycle shocks and the dynamic response of LFPR.² Second, 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 et al., 2017](#); [Amior and Manning, 2018](#); [Hornbeck and Moretti, 2018](#); [Yagan, 2019](#)) as well as the relationship between shocks

²[Balakrishnan et al. \(2015\)](#) also use a similar shift-share instruments, but for employment, while our paper uses output. Using the latter has many methodological advantages as we argue later on.

and migration (Howard, 2020; Monras, 2018; Cadena and Kovak, 2016). Our contribution to this literature is to provide estimates of the importance of compositional effects and to show how migration patterns differ by age. Third, we implement several of the recent methodological contributions in the shift-share empirical design (Adão et al., 2019; Borusyak et al., 2018; Goldsmith-Pinkham et al., 2020).

2 Research Design

We measure the cyclical nature 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 cyclical nature at the national level. In this section, 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 need a mean zero cyclical component, which rules out hysteresis by definition. Further, identification in those models relies on a trend component that is smooth and identifiable—a strong assumption for LFPR, given the sharp and changing nature of 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 :

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

where $y_{s,t}$ represents the labor market dependent variable of interest—e.g. 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 $X_{s,t}$ —which does not vary by group j —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. Nation-wide 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 the most common method used in past literature—vector autoregressions (VARs). [Stock and Watson \(2018\)](#) point out that VARs and local projections identify the same IRFs under standard conditions, but VARs may not correctly identify IRFs if the true IRFs are invertible. Additionally, [Olea and Plagborg-Møller \(2020\)](#) show that local projections have attractive properties for inference. 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 8.4](#), 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. GSP estimates are based on the factor incomes earned and other costs incurred in production (i.e., the GDI concept). For each state, the BEA 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 based around employment. Using GSP provides a measure of business cycle fluctuations at the state level that is more comprehensive than using employment only, which could omit productivity-driven business cycles. 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. Furthermore, using employment to measure business cycles is problematic when examining the response of employment as an outcome, since measurement error could create a mechanical relation even in the absence of an economic relation. Using output does not present this problem as it is a distinct economic variable and is measured by a separate data source from employment, as explained above. 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 percent shock to GSP growth.³ For $\beta^{(k)}$ to identify a causal effect of the GSP shock on $y_{s,t+k} - y_{s,t-1}$, it must be the case that, conditional on the set of controls, the growth rate of GSP in period t is uncorrelated with the error term:

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

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. The first-stage equation is as follows,

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

where

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

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.⁴ $\Delta GDP_{q,-s,t}$ represents the national gross domestic product growth in industry, q , for period, t , using the "leave-one-out" approach—that is, we calculate $GDP_{q,-s,t}$ by summing up $GSP_{q,s,t}$ across all states except for state s .

This formulation of the Bartik instrument relies on industry variation in output, rather than employment. Many previous studies, including [Blanchard and Katz \(1992\)](#), [Dao et al. \(2017\)](#), [Goldsmith-Pinkham et al. \(2020\)](#), and [Adão et al. \(2019\)](#), measure the response of employment to a Bartik 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.

³Assuming that the shock leaves GSP growth in other periods unaffected, this results in a permanent one percent shock to the level of GSP.

⁴During 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 $Bartik_{s,t}$ to be a valid instrument, it must meet the following conditions (Stock and Watson, 2018):

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

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

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

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

There are multiple interpretations of exogeneity for the Bartik instrument. The variation in the Bartik instrument comes from differential exposure to national shocks across regions, based on initial industry shares. Goldsmith-Pinkham et al. (2020) point out that this variation is equivalent to instrumenting with the industry shares directly, and therefore exogeneity of the instrument requires exogeneity of these shares. Borusyak et al. (2018) provide an alternative interpretation in which the national shocks are required to be exogenous. We are agnostic about which interpretation of the Bartik instrument provides exogeneity in this case and maintain the assumption that the Bartik term combining shares and shocks is exogenous.

2.2 Inference

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

Clustering: To quantify the uncertainty around our estimated impulse response functions, we compute heteroskedasticity-robust standard errors clustered at the state-level in our baseline specification. Adão et al. (2019) raise concerns that this approach may under-

state uncertainty in shift-share designs; however, our instrument is likely to be one for which state-clustered standard errors are appropriate.⁵ We validate this choice in [Section 8.3](#) 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 unemployment rates, LFPRs, and EPOPs 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 weights used in each regression.

2.3 Data

We combine state-level data from multiple sources to form an annual panel. In particular, labor market outcome variables consist of the unemployment rate, the labor force participation rate, and the employment-to-population ratio, which are all obtained from the Current Population Survey microdata. For each rate, we compute the average 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 additional results, we compute these rates for subgroups of the population.

In order to control for shifting demographics that are unrelated to the business cycle, we age-adjust each of our labor market outcome variables. That is, for an outcome $y_{i,s,t}$ for

⁵Our analysis is most similar to the results shown in Panel B of Table 6 in [Adão et al. \(2019\)](#), which shows that more sophisticated approaches to estimate confidence intervals are not meaningfully different from clustering by local labor market.

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 5.1](#).


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 Bartik instrument, 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

This section presents our main estimates of the cyclicity of the LFPR. We use our local projections regressions to measure the impulse response to a temporary shock to output growth among the CPS population, which includes non-institutionalized civilians ages 16 and over. [Figure 1](#) presents our estimates of the impulse response functions for the age-adjusted LFPR, unemployment rate, and 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 results show that the unemployment rate, LFPR, and EPOP all respond to cyclical shocks, but with varying timing. For the unemployment rate, a contractionary 1 percent shock to output causes a contemporaneous increase in the unemployment rate of 0.25 percentage points. The increase in the unemployment rate continues in the following year and peaks at 0.38 percentage points one year after the shock. Our estimate of the total increase in the unemployment rate due to a negative one percent shock to GSP is just below the range

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% 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.

of -0.5 to -0.4 of Okun's law coefficients estimated in the literature; see, for example, Ball et al. (2017). Following the peak one year after the shock, the unemployment rate steadily declines by about 0.1 percentage points per year until it returns to its pre-shock value about 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 et al. (2019).

The LFPR also shows a significant response to a negative GSP shock, but with a substantial delay compared to the unemployment rate. For example, the LFPR declines by less than 0.1 percentage 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 0.19 percentage points 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, several years after the unemployment rate has fully recovered.

The different patterns for the LFPR and unemployment rate reflect different cyclical pro-

files, which we can show formally with a nonlinear Wald-type test. We denote the coefficients tracing out the impulse response of the LFPR as $\{\beta_{\text{LFPR}}^{(k)}\}$ and the 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, i.e. $\beta_{\text{LFPR}}^{(k)} \equiv \phi_{\text{LFPR}} \cdot \beta_{\text{UR}}^{(k)}$ for each horizon k . Under the null, the ratio of coefficients $\frac{\beta_{\text{LFPR}}^{(k)}}{\beta_{\text{UR}}^{(k)}}$ is the same at every horizon k . To test this, we stack the samples used to estimate impulse responses for both variables and re-estimate Equation 1, from which we obtain a covariance matrix containing all coefficients for both impulse responses. 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 high cyclicity in the EPOP. The EPOP declines rapidly at the onset of the shock, reflecting the initial spike in the unemployment rate, and reaches its trough at -0.37 percentage point two years after the shock. Thereafter, the EPOP steadily recovers by about 0.05 percentage point per year until it is fully recovered eight 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 remaining weakness in the LFPR.


4 What Drives Long-Lived Cyclicity of Labor Force Participation?

Why does the LFPR take so long to respond and recover? One possible explanation for the long-lived cyclicity is that business cycle shocks lead people to make decisions which have persistent effects on their labor supply. Such decisions may include enrolling in schooling, staying at home and taking care for a family member, applying for disability benefits, or retiring.

To determine the extent to which each of these explanations may account for long-lived cyclicity, 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.⁶ Note that the “main reason for being out of the labor force” question is separate from the “want a job” question, and respondents have answers for both. 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 2. We show the IRFs only through eight years following the shock, since the estimates beyond lag eight become extremely noisy due to the limited sample.

Figure 2: 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% 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.

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.⁷ Initially, nonparticipants taking care of home/family constitute the largest response, with schooling close behind. However, nonparticipants reporting ill-

⁶For both of these questions, surveys before 1994 only asked these questions to about 1/4 of nonparticipants who are part of the Outgoing Rotation Groups in months 4 and 8 in sample.

⁷In 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.

ness 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 labor force participation. 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 response of labor force non-participants who say 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—as BLS’ U-5 measure does—are able to capture additional cyclical response beyond the main unemployment rate, but may fail to capture the long-lived cyclical response of participation.

5 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 of 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 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 25–39 year olds, likely due to out-migration. 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 cyclicality and hysteresis in response to local shocks.

5.1 Cyclicality of Adjusted and Unadjusted LFPRs

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

First, we estimate [Equation 1](#) using the unadjusted LFPR. [Figure 3](#) 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 individuals become persistently less likely to participate in the labor market as a result of the shock. However, in our case the persistently lower LFPR reflects changes in the state-level population distribution, and not a tendency for individuals to experience persistently lower participation conditional on their demographics. We find that people within a given age group should expect their labor force participation to fully recover after a shock, which is evidence against hysteresis.

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:


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

where $Y_{i,s,m,t}$ is a dummy 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 dummy 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 $X_{s,m,t}$ is a vector of state, month, year fixed effects.⁸ We include month-of-year dummy 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

⁸The age variables are single-year age dummies for ages 16 to 79 and a dummy variable for ages 80 years and older. The educational attainment dummies partition attainment into five categories: less than a high school degree, a high school degree, some college, a college degree, and more than a college degree. The race/ethnicity dummies partition the population into four groups: non-Hispanic white, non-Hispanic Black, Hispanic, and other. Marital status is a single dummy indicating whether an individual is married.

Figure 3: Cyclicity by Demographic Adjustment



Note: Each line shows the estimated coefficients from Equation 1 using the specified adjusted, unadjusted, and fitted-value LFPD 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% 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.

$\widehat{Y}_{i,s,m,t}^D$. With this fitted value, we calculate the demographically-adjusted LFPD 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 LFPD cyclicity. Figure 3 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 points to the age structure of the state-level population changing persistently in a way which would mechanically pull down the LFPD absent adjustment.

In Appendix Figure A.1 we repeat this exercise for the unemployment rate. In contrast to the LFPD, we find that demographics explain essentially none of the response of unem-

ployment, both immediately following the shock and in the long-run afterwards. This is likely a consequence of the fact that unemployment rates vary less over the life cycle than LFPRs, so changes in the age structure of the population affect less the unemployment rate.

5.2 Response of Population Composition to Cyclical Shocks

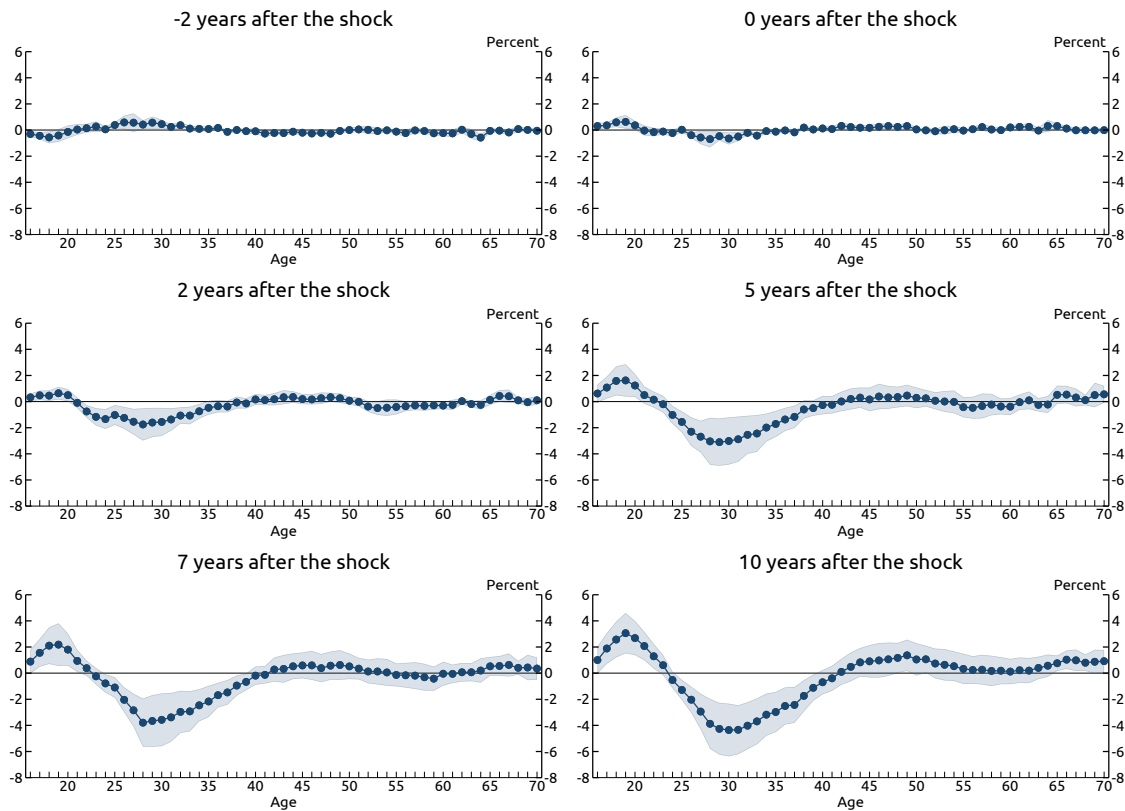
Why does the age-composition of the state-level population change in response to a business cycle shock? [Blanchard and Katz \(1992\)](#) provide empirical evidence that economic shocks at the state level trigger adjustments not only through unemployment, but also by triggering cross-state migration.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

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 cyclicalities may be related to local LFPR cyclicalities. 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 cyclicalities 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 cyclicalities, 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 the local projections [Equation 1](#) with the outcome $y_{s,t+k}$ being the log population of a single-year-age group in state s in period $t + k$. 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.⁹ Thus, our work in this section adds to the work by Blanchard and Katz (1992) and Dao et al. (2017) by estimating the population effects of single-year-age groups and identifying whether compositional effects—particularly, younger prime-age individuals—are driving the cyclical net migration results that those papers find.

Figure 4: 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 4, which shows the population response to a shock at -2, 0, 2, 5, 7, and 10 years

⁹We 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 domestic) to produce yearly population estimates for each age in each state. The covered-area population is slightly different from the civilian noninstitutional population, which is used to calculate LFPR and EPOP. The main difference is that the covered-area population includes active members of the armed forces and those in institutions (e.g. penal, mental facilities, and homes for the aged), whereas the civilian noninstitutional population does not include these groups. This distinction is not likely to matter in our analysis.

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 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, likely reflecting out-migration. Since 25 to 39 year olds are among the highest in LFPRs relative to other age groups, permanent declines in a state's population that are concentrated in this age range will also permanently lower its LFPR through compositional effects, all else equal.

There are several plausible reasons why the out-migration response might be concentrated in individuals ages 25 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 others states that are finishing school may be less likely to move to such a state. As a result, if a state experiences a recession, it could have a "missing generation" of recent graduates. This is consistent with the responses shown in Figure 4, 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.

6 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 individuals to more educated individuals, 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).

6.1 Age


The labor market performance of people between the ages of 25 and 54 (often referred to as prime-age people) is often used as a benchmark for the cyclical state of the labor market as a whole, as changes in demographics (such as the the aging of the baby boomers into their retirement years) may affect this group considerably less than the overall population. Although our main results control for these changes in demographics and thus give us a clean reading on the cyclical response of the labor market, understanding the cyclical response for the prime-age group is still of considerable interest, as 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.¹⁰

The cyclical response of the prime-age LFPR is similar to the overall response, albeit a bit smaller in magnitude. Figure 5 shows the estimated impulse response for the prime-age LFPR, along with unemployment rate and EPOP.¹¹ 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 points below its pre-shock level, before gradu-

¹⁰Although 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.

¹¹Unlike 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 5: 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% 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.


ally recovering and reaching its pre-shock level in year eight. The impulse responses for the unemployment rate and EPOP are also similar to our main results, although the standard errors for each of the prime-age responses are a bit wider than for the overall population.

Next we turn to other age ranges. Although the prime-age population often receives the most focus when discussing the health of the labor market, younger people (ages 16 to 24) and older people (ages 55 and over) comprise roughly half of the overall population and 30 percent of the labor force. Moreover, the shares of older people as a percentage of the population and labor force have slowly but steadily been increasing over the past 20 years, and are likely to continue increase into the near future, as the baby boomers continue to age in their retirement years. As a result, the responses of younger and older workers to a recession will have a large effect on how the labor market variables for the overall population respond an output shock.

The cyclicity of younger and older workers' labor market outcomes is especially difficult to measure from aggregate data. Decomposing changes in the LFPRs and EPOPs of younger and older people into changes caused by a recession as opposed to changes reflecting long-run trends is difficult, as the LFPR for both groups has trended sharply over the

past 30 years.¹² For younger workers, the 12 percentage point decline in their LFPR reflects both an increase in school enrollment rates and a decrease in the LFPR of those enrolled in school. For older workers, the 10 percentage point increase in their LFPR likely reflects a combination of changes in the age composition within the 55 and over group, an increase in the health capacity to work at older ages, and increases in the age at which individual’s can collect full retirement benefits through Social Security.¹³ Since these factors are likely to affect all states, our approach of leveraging business cycle shocks across states controls for these trends and allows us to identify the cyclical response of these groups.

Figure 6: Cyclical response 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% 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.

The LFPR for younger people reaches its trough sooner after a shock than does the overall LFPR, but its recovery is similarly delayed. The LFPR for young people reaches its trough by about one year after the shock—three years before the overall LFPR reaches its trough—at a level of 0.5 percentage points below its pre-shock value. Rather than beginning its recovery soon after reaching its trough, the LFPR for younger people lingers near at its trough level for an additional five years and does not begin its recovery until 7 years after the shock—2

¹²The sharp trends in EPOP for both younger and older people over the past 30 years entirely reflects their sharply trending LFPRs, as the unemployment rates for both groups show no clear, long-run trends.

¹³For general discussions on these factors, see Aaronson et al. (2014b), Montes (2018), and Bauer et al. (2019).

years after the overall LFPR starts its recovery. The point estimate of the LFPR of younger people never fully recovers, as it settles at about 0.2 percentage points below its pre-shock value, although the upper end of the confidence interval suggests we cannot rule out a full recovery. These results are consistent with our findings in 4, and a good portion of the long-lived cyclicalities of younger people may reflect them making persistent labor supply decisions around schooling.

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

6.2 Gender


Digging deeper into the prime-age LFPR responses, our results suggest that while both men and women have strong cyclicalities, the magnitudes and timing of their responses are quite different. For men, the initial point estimate response shown in the left panel of [Figure 7](#) 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 points 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 points 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 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.

6.3 Education

Labor market outcomes over at least the past 40 years have been quite different for lower- and higher-educated individuals. Indeed, the levels of the unemployment rates, LFPRs, and

Figure 7: Cyclicality 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% shock to GSP 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.

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

This divergence may, at least in part, be due to a long-term decline in the demand for lower-educated workers that is unrelated to the business cycle and caused, perhaps, by changes in technology and globalization; thus, to isolate cyclicalities 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 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 col-

Figure 8: Cyclicalities 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% shock to GSP 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.

lege degree experience essentially no variation in LFPR following a shock.¹⁴ 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.

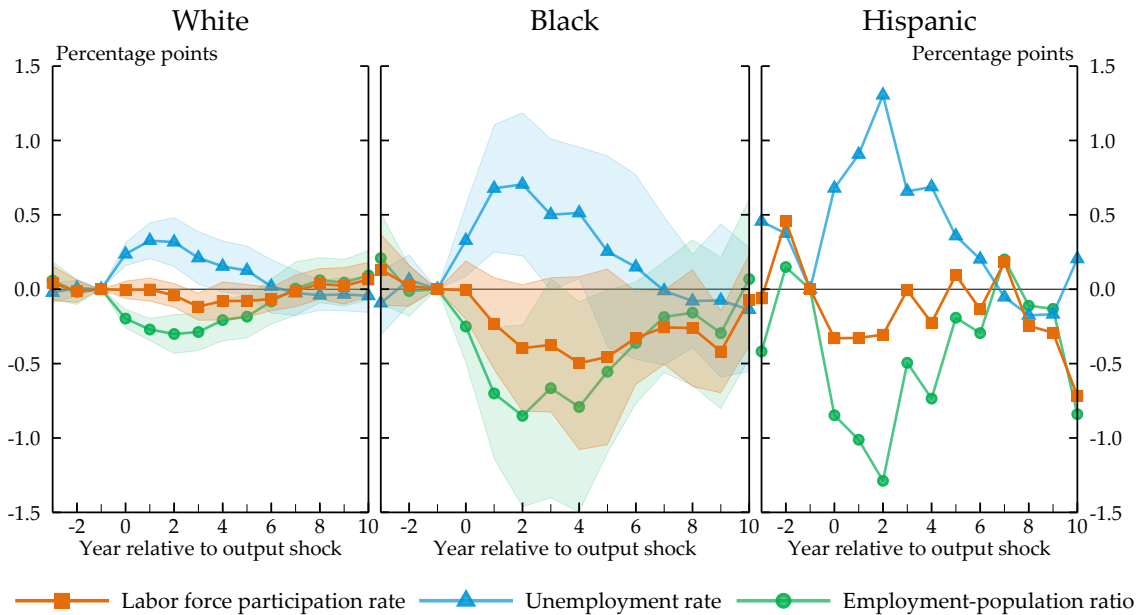
6.4 Race and Ethnicity

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 individuals in the CPS into racial and ethnic groups and estimate Equation 1 for each group, showing the results in Figure 9.

We find that shocks lead to larger and more long-lived declines in LFPR among minority groups. While the white LFPR falls by only 0.12 percentage point after a shock, the Black LFPR falls by 0.5 percentage point. The Black LFPR remains depressed for substantially

¹⁴We 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.

Figure 9: Cyclicity 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. “Black” includes people reporting both black and Hispanic, and “Hispanic” includes individuals reporting both white and Hispanic. Individuals not reporting either white, Black, or Hispanic are omitted. Coefficients are normalized to show the effect of a temporary -1% shock to GSP 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.

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.

7 What drives the shocks?

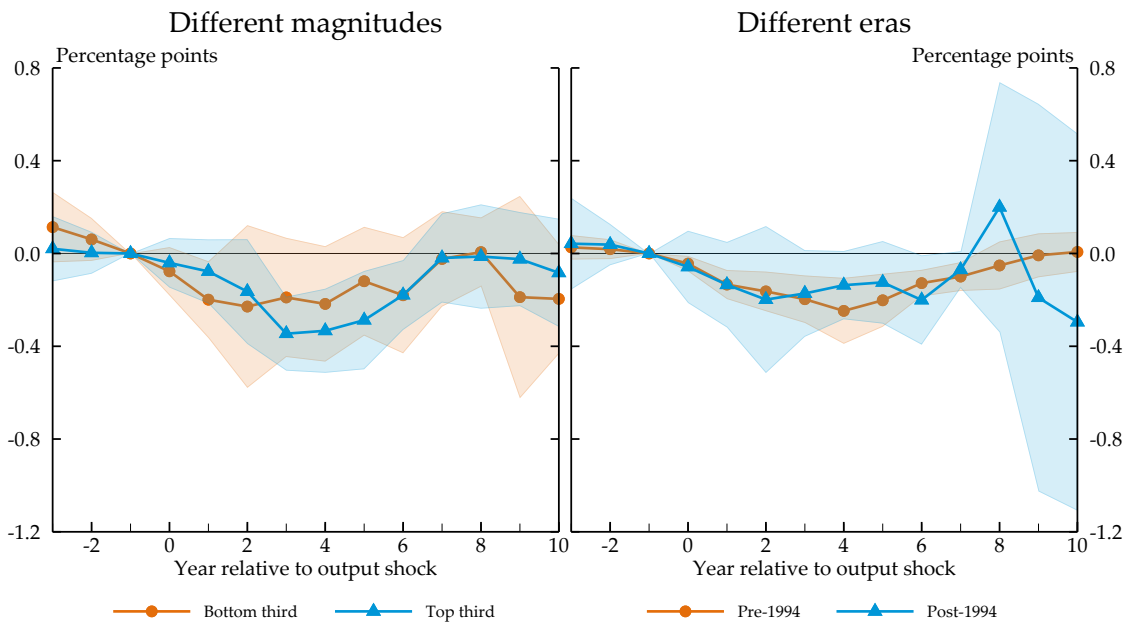
We examine what drives the variation in our output shock. We find similar responses to contractionary and expansionary shocks, suggesting that our effects are not being driven by asymmetries. More of our variation comes from the pre-1994 period, with estimates using only post-1994 data being similar overall but substantially noisier. The variation in the Bartik 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.

7.1 Expansions vs. contractions

Our instrument combines both expansionary and contractionary shocks to output, which could have different effects. If wages are downwardly rigid, as in Dupraz et al. (2019) and Murray (2019), a contractionary shock to output may lead to a greater decrease in labor force participation than an expansionary shock would raise it. 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.

Figure 10: 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% 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 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 10, we present the effects of expansionary shocks

(top third) and contractionary shocks (bottom third), normalizing both to show the effect of a -1 percent shock. Both impulse responses have similar patterns, and we cannot reject that the two are the same. This 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.

7.2 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.

To test whether the cyclicity 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 10. 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.

7.3 Decomposing the Bartik instrument

To further examine where the variation in our Bartik instrument comes from, we decompose the variation using the approach of Goldsmith-Pinkham et al. (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 compute the Rotemberg weights for each industry-year pair, we compute

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

where $\Delta GSP_{t,t-1}^{\perp}$ is GSP growth and $\Delta LFPR_{t+4,t-1}^{\perp}$ is the cumulative change in LFPR by four years after shock, both residualized on state and year fixed effects, Z'_{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

Table 1: Rotemberg weights in GSP Bartik 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 Bartik 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.

the leave-one-out growth rate, in order to align with the calculation of Rotemberg weights.¹⁵


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 Bartik instrument is that the variation comes from the industry shocks, as outlined in [Borusyak et al. \(2018\)](#).

Much of the variation in the Bartik 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 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 Bartik 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

¹⁵Our baseline results are little changed using the national growth rate instead of the leave-one-out growth rate.

Figure 11: 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% 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.

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.

7.4 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 Bartik instruments that are based purely on employment.

Figure 11 shows the estimated impulse response of productivity to a temporary -1 percent 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 11](#)), 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 productivity 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 Bartik instrument.

8 Robustness


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

8.1 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 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 Bartik instrument as the outcome variable. This impulse response is reported in the left panel of [Figure 12](#). 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.

Figure 12: Robustness Checks



Note: In the left panel, the line shows the estimated coefficients from Equation 1 using the Bartik 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% 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, and authors' calculations.

8.2 Controlling for house price growth

To verify the robustness of our results, we show that they remain unaffected when controlling for 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, e.g. through home price appreciation. Controlling for home price appreciation in the period before the shock addresses this concern. Panel I of Figure 12 shows that our estimates are little changed from the baseline if we control for home price growth.

8.3 Placebo

We cluster our standard errors at the state level in our baseline estimates, but Adão et al. (2019) point out that this may be insufficient in some circumstances. Our instrument exploits variation across places with different industry exposure, but places with a certain industry exposure may be subject to common shocks, giving rise to a particular structure for the variation attributable to unobserved shocks. 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 et al. \(2019\)](#). In place of our Bartik instrument, we estimate the reduced-form version of our main specification using a placebo Bartik 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 12](#). Unlike the cases examined by [Adão et al. \(2019\)](#), 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 suggests that our approach to inference is valid, and if anything is a bit conservative.

8.4 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 et al. \(2017\)](#) use VAR methods to 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 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 12](#) 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 cyclicity of LFPR with our approach.

9 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 cyclicity 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 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 of the LFPR poses problems for this view. In particular, we find that 5 to 7 years after a shock is a period in which the unemployment rate has essentially fully recovered, but the LFPR still has room to rise before it returns to its pre-shock level. During this period, observers who focus solely on the unemployment rate 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 of different demographic groups. Long-lived cyclical 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 cyclical in the LFPR is ignored.

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

References


- Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan**, “Declining labor force participation and its implications for unemployment and employment growth,” *Economic Perspectives*, 2014, 38 (4).
- Aaronson, Stephanie R., Mary C. Daly, William L. Wascher, and David W. Wilcox**, “Okun Revisited: Who Benefits Most from a Strong Economy?,” *Brookings Papers on Economic Activity*, 2019, 2019 (1), 333–404.
- Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher**, “Labor Force Participation: Recent Developments and Future Prospects,” *Brookings Papers on Economic Activity*, 2014, 45 (2 (Fall)), 197–275.
- Abraham, Katharine G. and Melissa S. Kearney**, “Explaining the Decline in the US Employment-to-Population Ratio: a Review of the Evidence,” *Journal of Economic Literature*, September 2020, 58 (3), 585–643.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales**, “Shift-Share Designs: Theory and Inference,” *The Quarterly Journal of Economics*, November 2019, 134 (4), 1949–2010.
- Amior, Michael and Alan Manning**, “The Persistence of Local Joblessness,” *American Economic Review*, July 2018, 108 (7), 1942–70.
- Balakrishnan, Ravi, Mai Dao, Juan Sole, and Jeremy Zook**, “Recent U.S. Labor Force Dynamics; Reversible or not?,” IMF Working Papers 15/76, International Monetary Fund April 2015.
- Ball, Laurence, Daniel Leigh, and Prakash Loungani**, “Okun’s Law: Fit at 50?,” *Journal of Money, Credit and Banking*, October 2017, 49 (7), 1413–1441.
- Bartik, Timothy J.**, *Who Benefits from State and Local Economic Development Policies?* number wbsle. In ‘Books from Upjohn Press.’, W.E. Upjohn Institute for Employment Research, December 1991.
- Bauer, Lauren, Emily Moss, Ryan Nunn, and Jay Shambaugh**, “Employment, Education, and the Time Use of American Youth,” Working Paper, The Hamilton Project, Brookings 2019.
- Blanchard, Olivier Jean and Lawrence F. Katz**, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992, 23 (1), 1–76.

- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” NBER Working Paper 24997 September 2018.
- Bound, John and Harry J Holzer**, “Demand shifts, population adjustments, and labor market outcomes during the 1980s,” *Journal of Labor Economics*, 2000, 18 (1), 20–54.
- Cadena, Brian C. and Brian K. Kovak**, “Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession,” *American Economic Journal: Applied Economics*, January 2016, 8 (1), 257–90.
- Cajner, Tomaz, Tyler Radler, David Ratner, and Ivan Vidangos**, “Racial gaps in labor market outcomes in the last four decades and over the business cycle,” 2017.
- Coglianesi, John**, “The Rise of In-and-Outs: Declining Labor Force Participation of Prime Age Men,” Working Paper 2018.
- Council of Economic Advisers**, “The Labor Force Participation Rate since 2007: Causes and Policy Implications,” mimeo 2014.
- Dao, Mai, Davide Furceri, and Prakash Loungani**, “Regional Labor Market Adjustment in the United States: Trend and Cycle,” *The Review of Economics and Statistics*, 2017, 99 (2), 243–257.
- Decressin, Jörg and Antonio Fatas**, “Regional labor market dynamics in Europe,” *European Economic Review*, 1995, 39 (9), 1627–1655.
- Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson**, “A Plucking Model of Business Cycles,” NBER Working Paper 26351 October 2019.
- Erceg, Christopher J. and Andrew T. Levin**, “Labor Force Participation and Monetary Policy in the Wake of the Great Recession,” *Journal of Money, Credit and Banking*, 2014, 46 (S2), 3–49.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, August 2020, 110 (8), 2586–2624.
- Hausman, J. A.**, “Specification Tests in Econometrics,” *Econometrica*, 1978, 46 (6), 1251–1271.
- Hornbeck, Richard and Enrico Moretti**, “Who benefits from productivity growth? Direct and indirect effects of local TFP growth on wages, rents, and inequality,” NBER Working Paper 24661 May 2018.

- Hornstein, Andreas and Marianna Kudlyak**, “Aggregate labor force participation and unemployment and demographic trends,” *Available at SSRN 3347310*, 2019.
- Howard, Greg**, “The Migration Accelerator: Labor Mobility, Housing, and Demand,” *American Economic Journal: Macroeconomics*, October 2020, 12 (4), 147–79.
- Jaimovich, Nir and Henry E. Siu**, “Job Polarization and Jobless Recoveries,” *The Review of Economics and Statistics*, 2020, 102 (1), 129–147.
- Monras, Joan**, “Economic shocks and internal migration,” 2018.
- Montes, Joshua**, “CBO’s Projection of Labor Force Participation Rates,” Working Papers 2018-04, Congressional Budget Office March 2018.
- Murray, Seth**, “Downward Nominal Wage Rigidity and Job Destruction,” Working Paper 2019.
- Olea, José Luis Montiel and Mikkel Plagborg-Møller**, “Local Projection Inference is Simpler and More Robust Than You Think,” Working Paper 2020.
- Shimer, Robert**, “Why is the U.S. Unemployment Rate So Much Lower?,” in “NBER Macroeconomics Annual 1998, volume 13,” MIT Press, January 1999, pp. 11–74.
- Staiger, Douglas and James H. Stock**, “Instrument Variables Regression with Weak Instruments,” *Econometrica*, 1997, 65 (3), 557–586.
- Stock, James H. and Mark W. Watson**, “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *The Economic Journal*, 2018, 128 (610), 917–948.
- Yagan, Danny**, “Employment Hysteresis from the Great Recession,” *Journal of Political Economy*, 2019, 127 (5), 2505–2558.
- Òscar Jordà**, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, March 2005, 95 (1), 161–182.

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% 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.