American Economic Journal: Applied Economics 2024, 16(3): 399–435
https://doi.org/10.1257/app.20220132

The Evolution of Local Labor Markets after Recessions†

By Brad Hershbein and Bryan A. Stuart*

This paper studies how US local labor markets respond to employment losses that occur during recessions. Following recessions from 1973 through 2009, we find areas that lose more jobs during the recession experience persistent relative declines in employment and population. Most importantly, these local labor markets also experience persistent decreases in the employment-population ratio, earnings per capita, and earnings per worker. Our results imply that limited population responses result in longer-lasting consequences for local labor markets than previously thought and that recessions are followed by persistent reallocation of employment across space. (*JEL* E32, J21, J31, J61, R23)

Recessions are a perennial feature of market economies. Since at least 1950, the US unemployment rate has tended to recover gradually after contractions (e.g., Dupraz, Nakamura, and Steinsson 2020; Hall and Kudlyak 2020), which raises the possibility that recessions have only modest long-run effects on the nationwide labor market. However, as recession severity can vary considerably across geographies, recessions could nonetheless have persistent consequences for *local* labor markets. The importance of understanding whether local areas also recover fully from recessions is underscored by a growing literature showing that local factors shape a range of outcomes—such as intergenerational mobility (Chetty and Hendren 2018a, b), health (Finkelstein, Gentzkow, and Williams 2021), and voting (Charles and Stephens 2013; Autor et al. 2020).

A series of influential studies suggest that local labor markets do recover completely from most recessions. The results in Blanchard and Katz (1992, hereafter BK) imply that, although employment losses persist, state employment-population ratios recover completely within ten years because of rapid population adjustments. Using additional years of data and a different source of identification to estimate the BK model, Dao, Furceri, and Loungani (2017) find that population is less responsive in the short run, but their estimates also imply full recovery of the employment-population ratio. Yagan

* Hershbein: W.E. Upjohn Institute for Employment Research (email: hershbein@upjohn.org). Stuart: Research Department, Federal Reserve Bank of Philadelphia (email: bryan.stuart@phil.frb.org). Camille Landais was coeditor for this article. We gratefully acknowledge funding from the 2018–2019 DOL Scholars Program (Contract # DOL-OPS-15-C-0060). We thank Steve Yesiltepe for excellent research assistance. For helpful feedback and suggestions, we thank the editor, referees, Tim Bartik, Leah Boustan, Katherine Eriksson, Andy Garin, Harry Holzer, Larry Katz, Lutz Kilian, Fabian Lange, Matt Notowidigdo, Chris Severen, Jay Shambaugh, Tara Sinclair, Anthony Yezer, and numerous seminar audiences. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper previously circulated as a working paper under the title "Recessions and Local Labor Market Hysteresis."
† Go to https://doi.org/10.1257/app.20220132 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

(2019) applies the BK methodology to study recessions and finds rapid recovery following the 1980–1982 and 1990–1991 recessions but slower recovery from the more severe Great Recession. Monras (2020) uses a different empirical strategy but also finds lasting effects on local areas after the Great Recession. One interpretation of this evidence is that recessions must be especially severe to generate persistent impacts on local labor markets. The accuracy of this interpretation has broad implications for our understanding of labor markets, economic opportunities available to workers and their children, and appropriate policy responses.

This paper studies the response of US local labor markets to employment losses that emerged during each recession between 1973 and 2009.[1] Specifically, we study how employment, population, and earnings evolve in local areas (metropolitan areas and commuting zones) where national recessions are more versus less severe. We draw upon multiple data sources, including those from the Bureau of Economic Analysis and Census Bureau, to create annual panels of longitudinally harmonized geographic areas stretching over five decades. We estimate regression models that relate the evolution of local economic activity to sudden employment changes that arise during recessions while controlling flexibly for changes in economic conditions at the regional level as well as prerecession population trends. This empirical strategy allows us to examine the extent to which local labor markets with larger employment losses during recessions recover relative to areas with smaller employment losses.

We find that declines in employment that emerge during recessions are extremely persistent. Across the five recessions that we study, a 10 percent decrease in metropolitan area employment during the recession—roughly the 90–10 percentile gap across areas for the Great Recession—on average leads to a 11 percent decrease in employment 7–9 years after the recession trough. The sudden decreases in employment that occur during recessions are not driven by differential pre-trends beforehand.

The consequences of these local employment declines depend on the extent of population adjustment. We find that metropolitan areas with larger employment losses experience population declines that begin during recessions and continue to grow for several years after the business cycle trough. The postrecession decrease in population is persistent but smaller than the decrease in employment. Due to this limited population response, local employment losses are followed by persistently lower employment-population ratios. On average, a 10 percent decrease in employment during a recession leads to a 5.6 percent (3.4 percentage point) decrease in the employment-population ratio. The change in the employment-population ratio accounts for about half of the decline in local area employment seven to nine years after the business cycle trough—with the decline in population explaining the remaining half. Moreover, these relative declines in employment-population ratios persist through at least 2019. Local employment losses during recessions are also followed by lasting decreases in earnings per capita and earnings per worker.

Our findings are consistent with local labor markets that experience larger employment losses during recessions facing a persistent downward shift in labor demand in the presence of a labor supply curve that is less than perfectly elastic but more elastic

[1] These recessions took place in 1973–1975, 1980–1982 (we pool the very short recession in 1980 with the longer one in 1981–1982), 1990–1991, 2001, and 2007–2009.

than population. Additional evidence suggests that our results reflect persistent consequences of labor market shifts that occur primarily during recessions, as opposed to a series of shifts taking place throughout the postrecession period. Consistent with this interpretation, we also show that our findings are not driven by secular changes in local economic activity that are correlated with local areas' prerecession industrial structure or demographic and labor market characteristics.

To further contextualize our results and corroborate our interpretation, we conduct several supplementary analyses. First, we find that relative declines in local employment are widespread across all sectors. Second, we use IRS data to show that the decline in population after the 2001 and 2007–2009 recessions arises from lower in-migration to local areas that experience larger employment losses. Out-migration actually falls after recessions in negatively affected areas. Third, we use individual-level data from the decennial census and the American Community Survey (ACS) to show that annual earnings declines tend to be more severe at the bottom and middle of the distribution. On average, about three quarters of the medium-term decline in annual earnings for those who remain employed arises from a reduction in hourly wages. Finally, using two complementary approaches, we present suggestive evidence that a change in the composition of residents due to selective migration does not account for most of the decline in local employment-population ratios or average earnings. Instead, the declines appear to stem mainly from lasting impacts on individuals, consistent with evidence on the effects of job displacement (e.g., Jacobson, LaLonde, and Sullivan 1993; Davis and von Wachter 2011; Lachowska, Mas, and Woodbury 2020; Schmieder, von Wachter, and Heining 2023).

Why do our results imply less recovery than the literature using vector autoregressions (VARs) (Blanchard and Katz 1992; Dao, Furceri, and Loungani 2017; Yagan 2019)? One potential explanation is that studies of local labor markets must rely on relatively short time series, which can lead to finite sample bias in VAR parameters. Using empirically relevant Monte Carlo simulations, we show that this finite sample bias leads VARs estimated in prior work to incorrectly imply convergence in cases where a decline in employment leads to permanent reductions in the employment-population ratio. The finite sample bias in our simulations would be of first-order importance even if researchers had access to 100 years of data. Moreover, we show that VAR estimates based on different years of state-level data or metro-level data imply complete recovery of the employment-population ratio, while event study regressions using state-level data are similar to our main results in suggesting more persistent declines. All of this evidence suggests that finite sample bias explains the difference in our results from those based on the BK VAR model.

The key contribution of this paper is evidence over a 50-year period on the response of local labor markets to employment losses that emerged during recessions. Our focus on recessions is motivated by two considerations. First, recessions have attracted substantial attention from researchers, policymakers, and the public. Second, as we show, recessions lead to sudden employment losses that break from preexisting trends, allowing us to generate transparent evidence on the evolution of local economic activity with flexible regression models. Our results show that local employment losses during recessions are followed by lasting shifts in the spatial distribution of employment and population. The results also show that relative reductions in employment-population ratios and earnings last longer than previously thought. Moreover, postrecession changes in local labor market outcomes are remarkably similar over the past five decades, which underscores the extent to which persistent local labor market disruption is a general feature of the US economy.

Our work complements recent research that uses local labor market variation to understand the consequences of recessions. Yagan (2019) uses tax data to provide evidence that people living in areas severely affected by the Great Recession experienced enduring employment and earnings losses. We differ from Yagan (2019) by focusing on how recessions affect local labor markets as opposed to individuals and by examining a larger number of recessions.[2] Monras (2020) provides empirical evidence that reduced in-migration accounts for essentially all of the population decline in areas hit harder by the Great Recession and develops a structural model to rationalize this fact. Our findings on in-migration are qualitatively similar. We differ from Monras (2020) in our empirical strategy and examination of more recessions and more outcomes.

Our work also complements several other studies that examine how local labor demand shifts, such as a change in manufacturing jobs, affect earnings, employment, population, and labor force participation (e.g., Bound and Holzer 2000; Freedman 2017; Amior and Manning 2018; Beaudry, Green, and Sand 2018; Gathmann, Helm, and Schönberg 2020; Notowidigdo 2020; Cajner, Coglianese, and Montes 2021; Garin and Rothbaum 2022). We provide new evidence by combining annual data—which directly reveal local labor market dynamics—and a research design that studies local employment shifts over a 50-year period. Additional evidence is particularly valuable because of the disagreement in the literature over whether shifts in local labor demand have persistent effects on wages and employment-population ratios and how, when, and why these relationships may have changed (Bartik 1993, 2015; Austin, Glaeser, and Summers 2018).[3]

Amior and Manning (2018) also show that incomplete adjustment of population to local employment shifts can generate persistent gaps in employment-population ratios. We differ in our use of sudden shifts in local employment that arise during recessions and our use of annual data—as compared to their analysis of predicted employment changes based on industrial structure using decadal data. Based on instrumental variable estimates of how employment responds to population and how population responds to employment, the model in Amior and Manning (2018) implies highly persistent labor demand innovations. In our setting, this would imply that areas that experience more severe recessions face additional negative labor demand shocks after recessions. We do not find evidence of such additional labor demand shocks, and our results are robust to controlling for prerecession local industry shares; both findings suggest that persistent local labor market declines do

[2] Rinz (2022) shows that the estimated effects in Yagan (2019) of the Great Recession on people's outcomes are smaller in magnitude when examining a broader range of ages, particularly because employment rates of individuals who were born from 1981–1996 recovered more quickly than other groups. Rinz (2022) also shows that the estimated impacts on people are slightly smaller than the impacts on places. Our results are similar to Rinz (2022) in finding that selective migration explains some but not most of the decline in local economic outcomes.

[3] Greenstone and Looney (2010) and Stuart (2022) provide evidence that recessions are followed by persistent declines in per capita earnings at the county level; our analysis goes considerably further by examining a larger range of outcomes, other levels of geography, additional recessions, and proximate mechanisms.

not arise in our setting because labor demand innovations are strongly correlated over time. Instead, our results suggest that the *effects* of specific labor demand disruptions that arise during recessions are persistent.

We emphasize that our finding of persistent local labor market declines is not inconsistent with aggregate economic recovery. The cross-sectional identifying variation we use identifies relative differences in the evolution of local labor market outcomes between areas that experience more or less severe employment losses during recessions.[4] A persistent relative decline does not imply that an area fails to recover in an absolute sense but rather that a gap remains between that area and one that experienced a less severe recession. These relative differences most directly shed light on the distributional consequences of recessions and the efficiency costs associated with incomplete local labor market adjustments.

I. Conceptual Framework

To guide our empirical analysis, we draw on a simple conceptual framework and previous research to describe how local labor markets might evolve after recessions. This discussion informs our empirical strategy and the interpretation of our results.

Consider a local labor market that experiences a decline in employment during a recession. Over a short horizon of two to three years, the most natural catalyst of this fall in employment is a downwards shift in labor demand. The fall in demand could stem from many possible sources, such as an increase in interest rates or oil prices or a decline in consumer sentiment. Employment will fall during a recession if labor supply is not perfectly inelastic in the short run. Wages will fall if labor supply is less than perfectly elastic, and the employment-population ratio will also fall if the labor supply elasticity is larger than the population elasticity.[5]

After the recession, the local labor market could recover to varying degrees. The extent of recovery can be summarized by focusing on two questions.

First, are there lasting declines in employment and population? On the one hand, the local labor market could exhibit complete recovery in terms of these variables if the downward shift in labor demand is temporary and there is no shift in labor supply. For example, this pattern would arise if firms temporarily laid off workers or reduced their hours and there was no change in the nonwage determinants of labor supply and population, such as quality of life. On the other hand, the local labor market could experience a persistent decline in employment or population if the negative shift in labor demand persists and the supply of employment or population is not perfectly inelastic.

Second, are there lasting declines in the employment-population ratio and wages? This question is mainly of interest in the case where there are lasting declines in

[4] Other papers studying local labor markets also identify relative differences (e.g., Blanchard and Katz 1992; Autor, Dorn, and Hanson 2013; Amior and Manning 2018).

[5] There are several possible explanations for why the labor supply elasticity could exceed the population elasticity at any horizon. For example, working age individuals might be more mobile than other individuals (such as retirees). Working age individuals could also care more about employment opportunities than individuals that are not working. Finally, individuals might adjust their labor supply without moving, possibly by dropping out of the labor force when wages fall below their reservation level.

employment and population. If individuals' labor supply and migration choices are extremely sensitive to local job opportunities, then a combination of labor force exits, higher out-migration, and lower in-migration could reequilibrate the local labor market at near its original employment-population ratio and wage level. In contrast, if the supplies of labor and population to a local area are both relatively inelastic, then there could be lasting declines in wages. If labor supply is more responsive than population, then the employment-population ratio would remain depressed as well.[6]

Existing research tends to find evidence of lasting declines in employment and population in response to shifts in labor demand. One set of papers has found these results by estimating VARs (Blanchard and Katz 1992; Dao, Furceri, and Loungani 2017; Yagan 2019).[7] Other papers have studied the consequences of employment changes predicted by the interaction of preexisting industrial structure and nationwide industrial shifts, often over decadal intervals (Bartik 1991; Bound and Holzer 2000; Amior and Manning 2018; Notowidigdo 2020). The evidence in Monras (2020) from the Great Recession also indicates a persistent decline in population.

There is less agreement about whether shifts in labor demand are followed by persistent changes in the employment-population ratio and wages. Estimates from VARs generally imply that the employment-population ratio, unemployment rate, labor force participation rate, and wages recover fully within a decade (Blanchard and Katz 1992; Dao, Furceri, and Loungani 2017; Yagan 2019).[8] On the other hand, papers that study employment changes predicted by industrial structure tend to find evidence of lasting changes in wages and the employment-population ratio (Bartik 1991; Bound and Holzer 2000; Amior and Manning 2018; Notowidigdo 2020).[9]

Assessing the extent of declines in the employment-population ratio and wages is particularly important because these measures more directly reflect the economic opportunities available to the average person in an area. In turn, these results inform our understanding of fundamental features of local labor markets and appropriate policy responses to labor demand shifts. We next describe our approach to studying these questions. After presenting our main results, we describe in detail how our findings relate to these cases and past work.

[6] It is reasonable to expect that the labor supply elasticity is larger than the population elasticity (see footnote 5). However, if the labor supply elasticity is smaller than the population elasticity, then a persistent decline in local labor demand could lead to an increase in the employment-population ratio.

[7] These papers differ in the terms of magnitudes: BK estimate that employment recovers by almost 40 percent after the shock; Dao, Furceri, and Loungani (2017) estimate recovery between 20–40 percent; and Yagan (2019) estimates even more recovery following the 1980–1982 and 1990–1991 recessions. The evidence in Yagan (2019) comes from estimating the BK VAR using data from 1978–2007 and then calculating averages for states where the VAR-implied recession shock was more or less severe. The presence of pre-trends in this simple model makes it difficult to conclusively say how persistent the employment changes are in these results.

[8] The estimates in Yagan (2019) imply similarly rapid recovery following the 1980–1982 and 1990–1991 recessions but slower recovery after the Great Recession.

[9] Monras (2020) also finds evidence of lasting impacts on wages after the Great Recession.

II. Data and Empirical Strategy

A. *Data*

We compile several public-use datasets to measure local economic activity. These datasets are constructed by government agencies using administrative data. Employment is available from the Bureau of Economic Analysis Regional Economic Accounts (BEAR) (Bureau of Economic Analysis 1969–2019*b*), Census County Business Patterns (CBP) (Census Bureau 1970–1994, 1995–2017), and Quarterly Census of Employment and Wages (QCEW) (Bureau of Labor Statistics 1975–2019*d*).[10] BEAR and CBP data are available starting in 1969, while QCEW data are available from 1975 onward. BEAR data also contain aggregate earnings.[11] We use the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) (National Cancer Institute 1969–2019) data for annual population estimates, which are available by sex, race, and age. To measure in- and out-migration, we use the Internal Revenue Service Statistics of Income (SOI) (Internal Revenue Service 1993–2019) data.[12] Finally, we use tabulations and microdata from the decennial census and ACS to examine the earnings distribution and composition changes.[13]

With the exceptions of the decennial census and ACS microdata, all of the datasets are available at the county level. The census and ACS are available at the Public Use Microdata Area level, which we map to other geographies using crosswalks available from the Geocorr program of the Missouri Census Data Center. Consequently, we can examine changes in economic activity for metropolitan areas and commuting zones.[14] Both types of areas are composed of counties, so it is straightforward to map our county-level data into metropolitan areas or commuting zones. A slight complication is that definitions of metropolitan areas and commuting zones change over time; we use Core Based Statistical Areas (CBSAs) as defined by the Office of Management and Budget in December 2003 and commuting zones as defined by US Department of Agriculture and based on the 2000 census. Although we focus on metropolitan areas because of their greater size and thicker labor markets, we show

[10] Because employment counts are often suppressed for small counties and industries in CBP data, we adopt the imputation procedure of Holmes and Stevens (2002) when necessary. Details are in the online Data Appendix. Results from this approach agree closely with WholeData, which uses a linear programming algorithm to recover suppressed employment estimates (Bartik et al. 2019).

[11] More specifically, BEAR data contain earnings by both place of residence and place of work. Since wage and salary employment is available only by place of work, we use earnings by place of work. We define earnings as wages, salaries, and supplements (benefits), and we adjust earnings for inflation using the personal consumption expenditures deflator in 2019 dollars (Bureau of Economic Analysis 1969–2019*a*). As discussed below, our results are similar when alternatively measuring earnings by place of residence.

[12] SOI data are available starting in the 1990s. Although they capture moves only for tax filers, SOI data are considered a high-quality source for point-to-point migration flows and have been used in several papers (e.g., Kaplan and Schulhofer-Wohl 2012, 2017; Wilson 2022).

[13] We use versions of these tabular data and microdata from NHGIS and IPUMS, respectively (Manson et al. 2019; Ruggles et al. 2019). The online Data Appendix describes the processing of these data and how we link individuals to our geographies of interest.

[14] Metropolitan statistical areas are defined by the Office of Management and Budget as having "at least one urbanized area of 50,000 or more population, plus adjacent territory that has a high degree of social and economic integration with the core as measured by commuting ties" (Office of Management and Budget 2003). Commuting zones are defined based on commuting patterns and do not have a minimum population threshold or urban requirement (Tolbert and Sizer 1996). We do not examine counties because these are often too small to constitute local labor markets, our area of focus.

that our main results are robust to using commuting zones, which, unlike metropolitan areas, cover the entire United States.[15]

B. *Empirical Strategy*

Our empirical strategy relies on cross-sectional variation in sudden employment changes that occur during nationwide recessions. We use this variation to estimate how the postrecession evolution of local labor market outcomes varies with the severity of recessions.

Our preferred approach is to stack recessions together and estimate the following regression:

$$y_{i,t} - y_{i,p(r)-2} = \sum_{\tau=p(r)-4}^{p(r)+12} s_i^r \mathbf{1}\{t = \tau\} \delta_\tau + \mathbf{x}_i^r \beta_t^r + \varepsilon_{i,t}^r, \tag{1}$$

where $y_{i,t}$ is a measure of local economic activity in location i and year t; s_i^r is the severity of recession r, measured as the log employment change in location i from the nationwide business cycle peak to trough (multiplied by -1); $\mathbf{1}\{t = \tau\}$ is an indicator for year t being equal to τ, which ranges from 4 years before to 12 years after the nationwide recession start year $p(r)$; $\mathbf{x}_i^r$ is a vector of recession-specific, time-invariant control variables; and $\varepsilon_{i,t}^r$ is an error term. The term $y_{i,p(r)-2}$ is the outcome measure in location i two years before the nationwide recession start so that the left-hand side of equation (1) is the within-location change in the outcome relative to a fixed, prerecession period.

The key parameter of interest, δ_τ, describes the relationship between the change in employment during the recession and the change in local economic activity as of year τ relative to the nationwide recession start. Because the left-hand side of equation (1) is a within-location change, this approach controls for time-invariant cross-sectional differences. We normalize the δ_τ coefficient to equal zero two years before the recession start (i.e., $\delta_{p(r)-2} = 0$) to facilitate comparisons across recessions. We choose two years before the recession start as the normalization year because the exact timing of recessions is uncertain and there is variation in when aggregate economic indicators decline.[16] The δ_τ parameters vary freely across years relative to the recession start, which is useful for identifying empirical patterns without imposing possibly incorrect constraints. Moreover, stacking the five recessions into a single regression in event time allows us to increase precision and focus on central tendencies.[17] This reduced-form approach can capture a wide variety of demand and supply adjustments after the recession.

[15] Metropolitan areas, consistently defined, cover 80–90 percent of people and jobs throughout our sample, with this share growing over time.

[16] Because we show the entire range of estimates of δ_τ, it is straightforward to see how our estimates would change with a different normalization year.

[17] Cengiz et al. (2019) adopt a similar stacked event study approach in their analysis of the minimum wage. We present results for each recession separately in the online Appendix, as referenced below.

We measure local recession severity using annual employment data from BEAR. We modify National Bureau of Economic Research (NBER) business cycle peak and trough dates to account for our use of annual data. Specifically, we construct s_i^r using the log employment change for each geography between 1973 and 1975, 1979 and 1982, 1989 and 1991, 2000 and 2002, and 2007 and 2009.[18] We use wage and salary employment (private and public) to measure recession severity, as coverage of the self-employed is incomplete and varies over time. Variation across areas in employment losses during recessions can arise from differences in industrial specialization (e.g., recessions could decrease demand for automobiles) or even finer differences in the products that are made in an area (e.g., recessions could particularly decrease demand for more expensive trucks and SUVs). Idiosyncratic shocks to a single large firm could also generate local employment losses (cf. Gabaix 2011; Salgado, Guvenen, and Bloom 2019).

Estimates of δ_τ can be interpreted as isolating the differential response of local economic outcomes with respect to recession severity if s_i^r is exogenous to changes in residual determinants of local labor market outcomes, $\varepsilon_{i,t}^r$, conditional on the controls in the regression. In addition to controlling for time-invariant differences across local areas, we include several variables in $\mathbf{x}_i^r$ to bolster the validity of this interpretation. First, we include census division fixed effects to flexibly capture broad shifts in local labor demand and supply that are not driven by recessions, such as the rise of the sunbelt (Glaeser and Tobio 2008). Second, we include *pre*recession population growth to adjust for secular shifts in local labor supply.[19] The coefficient vector on these controls, β_t^r, varies freely across years and recessions for increased flexibility. In unreported results, we find that estimates are very similar when additionally controlling for prerecession employment growth with coefficients that vary by year and recession. Estimates of δ_τ for prerecession years allow us to directly examine the presence of pre-trends, and estimates of δ_τ for postrecession years shed light on whether areas that experience larger employment losses during recessions are differentially exposed to nonrecession economic shocks (which would show up as subsequent spikes or jumps in δ_τ). We cluster standard errors at the metro level to allow for arbitrary autocorrelation in the error term $\varepsilon_{i,t}^r$ across years and recessions.

There are several notable aspects to our analysis of how local labor markets evolve after employment changes that occur during recessions. First, recessions feature both general declines in economic conditions and increased dispersion in economic conditions across different areas (e.g., Dao, Furceri, and Loungani 2017). While many metro areas experience absolute job losses during recessions, online Appendix Figures 1 and 2 show that several areas also see gains in employment

[18] The NBER recession dates are November 1973 to March 1975, January to July 1980, July 1981 to November 1982, July 1990 to March 1991, March to November 2001, and December 2007 to June 2009.

[19] Controlling for baseline levels or pre-trends of economic outcomes is common (e.g., Autor, Dorn, and Hanson 2013; Dix-Carneiro and Kovak 2017; Hershbein and Kahn 2018). Given the challenge of controlling directly for all relevant local labor supply shifters (e.g., due to a wide range of natural and cultural amenities), we opt to control for prerecession population growth. We control for the log change in population for ages 0–14, 15–39, 40–64, and 65 and above. We construct these population variables using SEER data, which are available starting in 1969. The prerecession population growth years are 1969–1973 (for the 1973–1975 recession), 1969–1979 (for the 1980–1982 recession), 1979–1989 (for the 1990–1992 recession), 1990–2000 (for the 2001 recession), and 1997–2007 (for the 2007–2009 recession).

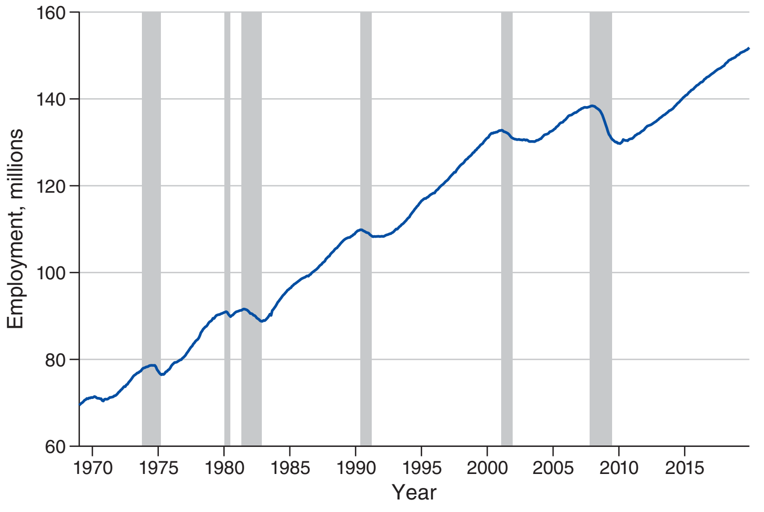

Figure 1. Aggregate Employment and Recessions, 1969–2019

Notes: Figure shows seasonally adjusted national nonfarm employment. The shading indicates NBER national recession dates.

Source: Authors' calculations using Bureau of Labor Statistics Current Employment Statistics.

over the two-to-three-year recession horizons we examine. Second, as mentioned above, our use of the actual log employment change as the key explanatory variable of interest implies that our regressions partly reflect idiosyncratic shifts in local labor demand. It is not clear a priori whether shifts in employment that include these idiosyncratic factors would lead to more or less persistent changes in local economic conditions than shifts in employment that are predicted by industry shift-share shocks (Bartik 1991). On the one hand, employment shifts that include idiosyncratic shocks might better capture the consequences of job losses at important establishments or plant closures. On the other hand, shift-share shocks might better reflect structural changes in the economy.[20] Third, the consequences of a change in actual employment or a shift-share shock could differ across recessions because of heterogeneity in the macroeconomic shock or the areas that are exposed to the shock (e.g., Adão, Kolesár, and Morales 2019). These considerations motivate our approach of estimating impacts separately for each recession for transparency, estimating stacked regressions to increase precision and focus on central tendencies, and comparing results that rely on variation in the actual employment change in an area to those that rely on variation from the predicted employment change based on a shift-share shock.

[20] That said, shift-share shocks based on the two-to-three-year recession horizons we study might be less likely to capture structural changes like the decline in manufacturing than the ten-year horizons that are used in other work (e.g., Bound and Holzer 2000; Goldsmith-Pinkham, Sorkin, and Swift 2020).

C. *The Severity of Recessions across Time and Space*

Before moving to estimates of equation (1), we describe the characteristics of the five recessions that are our focus. Figure 1 displays aggregate seasonally adjusted nonfarm employment from the Current Employment Statistics from 1969 to 2019 (Bureau of Labor Statistics 1969–2019). Nationwide employment more than doubled over this period. This growth was interrupted by five recessions (combining the two in the early 1980s), as indicated by the vertical shaded bars in the graph. While there is little consensus on the macroeconomic causes of each recession, the drivers almost certainly differ (Temin 1998). The 1973–1975 and 1980–1982 recessions followed increases in the price of oil and subsequent increases in interest rates by the Federal Reserve. There is less agreement on the causes of the 1990–1991 recession (Temin 1998). The 2001 recession followed the dot-com bubble, and the 2007–2009 recession followed tumult in housing and financial markets.

Using annual data from BEAR, Table 1 shows the national changes in employment from business cycle peak to trough for each recession, both overall and for major industrial sectors. The recessions vary in overall magnitude, from a 3 percent employment decline during the Great Recession to a 1 percent increase from 1989 to 1991, with the others falling in between. The bottom right panel of the table reports pooled employment changes across recessions. The results show that manufacturing and construction usually experience the largest proportional employment decline. The patterns of employment changes for other industries differ more across recessions.

Figure 2 displays the frequency with which each area experienced a severe recession over the sample horizon. We define a metropolitan area as having a severe recession if it experienced a log employment change worse than the median area for a given recession. The Detroit and Chicago metros, for example, experienced downturns worse than the median for all five recessions, while the Houston metro did so only in 2001. The distribution in severity frequency is roughly symmetric, with a similar number of metros experiencing zero or one severe recessions (112) as those experiencing four or five (105). As a result, there is considerable variation across recessions in whether a given area faces a severe employment loss.[21]

We show the serial correlation in recession severity in Table 2. Panel A shows the raw correlations across metros in log employment changes for each pair of recessions. As suggested by Figure 2, the serial correlation is positive but moderate. Consistent with the different origins of the recessions as well as temporal changes in industrial mix, the pattern is not monotonic across time. We also show in panel B the correlations after partialing out fixed effects for the nine census divisions and in panel C the correlations after additionally controlling for prerecession population growth. These controls tend to slightly reduce the magnitudes of the correlations, but positive serial correlation remains in a few cases. The regression estimates below suggest that serial correlation in recession severity has relatively little impact on our results. We also

[21] This result is also apparent when examining log employment changes separately for each recession (online Appendix Figure 1). Moreover, a substantial share of areas see absolute increases in employment growth during each episode (online Appendix Figure 2).

Table 1—Aggregate Employment Changes, by Recession

	1973–1975 Recession			1980–1982 Recession			1990–1991 Recession		
	Share of peak year employment (1)	Log employment change (2)	Employment change (3)	Share of peak year employment (4)	Log employment change (5)	Employment change (6)	Share of peak year employment (7)	Log employment change (8)	Employment change (9)
Total	1.000	0.004	421,100	1.000	0.010	1,123,200	1.000	0.011	1,531,000
Manufacturing	0.216	−0.090	−1,758,600	0.196	−0.110	−2,230,100	0.150	−0.049	−962,800
Services	0.203	0.053	1,041,400	0.220	0.103	2,606,900	0.276	0.060	2,264,500
Government	0.177	0.046	792,000	0.168	0.008	149,000	0.156	0.023	493,000
Retail trade	0.159	0.010	153,300	0.161	0.020	359,600	0.168	0.005	110,800
Finance, insurance, real estate	0.076	0.027	192,700	0.079	0.037	322,200	0.080	−0.014	−146,000
Transportation and public utilities	0.054	−0.018	−91,400	0.052	0.003	17,400	0.048	0.034	220,600
Construction	0.054	−0.084	−410,000	0.054	−0.096	−536,900	0.054	−0.065	−451,500
Wholesale trade	0.048	0.073	341,800	0.052	0.008	44,900	0.050	−0.012	−76,200
Mining	0.008	0.140	114,100	0.011	0.264	350,800	0.008	−0.025	−26,000
Agriculture, forestry, fisheries	0.006	0.073	45,800	0.008	0.043	39,400	0.010	0.077	104,600

	2001 Recession			2007–2009 Recession			Pooled log employment change	
							Mean	SD
Total	1.000	−0.000	−62,700	1.000	−0.034	−5,866,000	−0.002	0.017
Manufacturing	0.109	−0.120	−2,004,900	0.082	−0.147	−1,982,600	−0.103	0.033
Services	0.409	0.022	1,504,500	0.432	−0.012	−886,900	0.045	0.038
Government	0.141	0.027	638,000	0.137	0.018	452,000	0.025	0.013
Retail trade	0.114	−0.015	−268,300	0.107	−0.064	−1,171,600	−0.009	0.030
Finance, insurance, real estate	0.082	0.019	260,100	0.094	0.025	426,900	0.019	0.017
Construction	0.059	0.013	128,500	0.064	−0.190	−1,975,100	−0.084	0.065
Transportation and public utilities	0.038	−0.022	−133,000	0.037	−0.061	−385,500	−0.013	0.031
Wholesale trade	0.039	−0.027	−169,900	0.037	−0.070	−443,300	−0.006	0.047
Mining	0.005	−0.012	−9,000	0.006	0.107	114,300	0.095	0.106
Agriculture, forestry, fisheries	0.005	−0.010	−8,700	0.005	−0.017	−14,200	0.033	0.040

Notes: Table reports nationwide wage and salary employment changes during recessions. Employment changes are from 1973–1975, 1979–1982, 1989–1991, 2000–2002, and 2007–2009. The 1973–1991 data are based on Standard Industrial Classification (SIC) industries, and the 2000–2009 data are based on North American Industry Classification System (NAICS) industries. Industry changes may not sum to total changes due to rounding. The bottom right panel shows means and standard deviations of log employment changes across the five recessions.

Source: Authors' calculations using BEAR data.

control for the severity of previous recessions as an additional robustness check and show that these controls do not appreciably change the results.

Table 3 describes the characteristics of metropolitan areas that experience a more versus less severe recession (defined as whether the log employment change is above

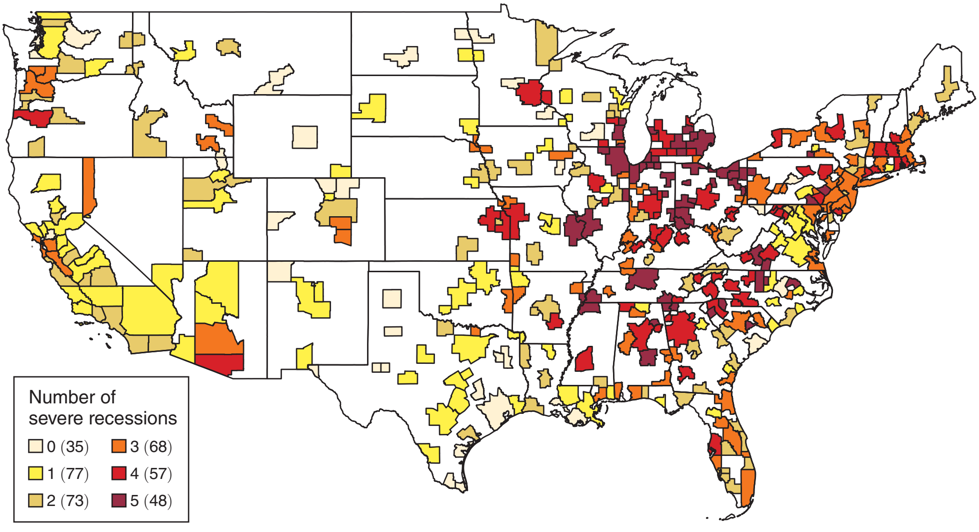

Figure 2. Frequency of Severe Recessions, by Metropolitan Area, from 1973–2009

Note: We define an area as experiencing a severe recession if its log employment change for a given recession is less than the median across the 358 metropolitan areas for that recession.

Source: Authors' calculations from BEAR.

or below the median). We measure these characteristics using the closest decennial census to the recession start year, except for the 2007–2009 recession, which is measured using the 2005–2009 ACS. The largest consistent difference between areas that experience a more versus less severe recession is the manufacturing employment share, though this difference has decreased considerably over time. The other differences vary across recessions and are generally small.[22] The variables in Table 3 include both sources of recession severity and factors that might influence the response of local areas to decreases in employment. We directly examine changes in some of these variables while also examining changes in worker composition to better understand related mechanisms.[23]

[22] Online Appendix Table 1 reports *p*-values of the differences in these characteristics between areas facing more versus less severe recessions. Additionally, online Appendix Tables 2 and 3 show descriptive estimates of the relationship between prerecession metro area characteristics and, respectively, the metro-level log employment change during recessions (our primary regressor of interest) and the shift-share-predicted log employment change. The manufacturing employment share is strongly and consistently predictive of more severe downturns in both cases, but other characteristics are more variable across measures and recessions. As indicated by the R-squareds from the tables, the log employment change during a recession captures more idiosyncratic factors than does the predicted shift-share instrument.

[23] We examined whether postrecession changes in economic activity varied with prerecession levels of these variables but found little evidence of such heterogeneity.

Table 2—Correlation of Metropolitan Area Recession Severity

	Change in log employment during recession years				
	1973–1975	1979–1982	1989–1991	2000–2002	2007–2009
Panel A. Unadjusted					
1973–1975	1.000				
1979–1982	0.386	1.000			
1989–1991	0.459	0.154	1.000		
2000–2002	0.446	0.412	0.281	1.000	
2007–2009	0.354	0.210	0.002	0.155	1.000
Panel B. Adjusted for census division					
1973–1975	1.000				
1979–1982	0.327	1.000			
1989–1991	0.275	0.170	1.000		
2000–2002	0.291	0.304	0.234	1.000	
2007–2009	0.363	0.071	−0.044	0.091	1.000
Panel C. Adjusted for census division and prerecession population growth					
1973–1975	1.000				
1979–1982	0.258	1.000			
1990–1991	0.161	0.018	1.000		
2000–2002	0.144	0.084	0.098	1.000	
2007–2009	0.400	0.279	0.050	0.212	1.000

Notes: Table reports correlations of log wage and salary employment changes across recessions for 358 metropolitan areas. Panel B reports correlations after partialling out census division fixed effects, and panel C partials out census division fixed effects and prerecession population growth.

Source: Authors' calculations using BEAR data.

III. The Postrecession Evolution of Local Economic Activity

A. *Employment*

We begin with estimates of equation (1) for log employment in metropolitan areas. Panel A of Figure 3 presents estimates from the stacked regression.[24] We include four years before the start of each recession to capture any pre-trends, and we follow areas for 12 years afterwards. Specification 1, shown in red (circles), includes only census division fixed effects in $\mathbf{x}_i^r$. Our preferred specification 2 (solid blue line) also controls for prerecession population growth for ages 0–14, 15–39, 40–64, and 65 and above. Specification 3 (green squares) adds the severity of the previous recession, which is possible for all but the 1973–1975 recession. Finally, specification 4 (orange triangles) further includes the severity of *all* previous recessions since 1973. In all cases, we allow the coefficient vector β_t^r to vary freely across years and recessions (e.g., we interact division fixed effects with year-by-recession fixed effects).

Overall, there is some weak evidence of negative pre-trends from specification 1, indicating that employment was gradually declining beforehand in areas

[24] Recession-specific estimates are in online Appendix Figure 3.

Table 3—Characteristics of Metropolitan Areas with More versus Less Severe Recessions

	Recession					
	1973–1975		1980–1982		1990–1991	
Prerecession characteristic	Less severe	More severe	Less severe	More severe	Less severe	More severe
Manufacturing emp. share	0.141	0.254	0.140	0.236	0.131	0.179
Mining emp. share	0.013	0.004	0.013	0.005	0.013	0.005
Construction emp. share	0.052	0.051	0.058	0.051	0.055	0.053
Finance, insurance, real estate emp. share	0.062	0.059	0.073	0.063	0.068	0.065
Population (1,000s)	333.1	595.4	552.9	430.6	329.8	768.2
Log population growth	0.090	0.066	0.247	0.108	0.137	0.078
Employment-population ratio	0.518	0.537	0.534	0.547	0.546	0.579
Real earnings per capita (1,000s)	19.7	21.0	21.5	23.2	23.5	26.4
Share with BA degree or more	0.120	0.096	0.172	0.142	0.195	0.183
Non-White share	0.145	0.133	0.209	0.122	0.189	0.188
Foreign-born share	0.029	0.027	0.048	0.028	0.045	0.043

	Recession			
	2001		2007–2009	
Prerecession characteristic	Less severe	More severe	Less severe	More severe
Manufacturing emp. share	0.096	0.163	0.082	0.110
Mining emp. share	0.008	0.003	0.008	0.002
Construction emp. share	0.059	0.056	0.060	0.067
Finance, insurance, real estate emp. share	0.066	0.064	0.073	0.079
Population (1,000s)	531.6	732.4	618.7	744.7
Log population growth	0.162	0.096	0.091	0.117
Employment-population ratio	0.591	0.632	0.612	0.585
Real earnings per capita (1,000s)	28.3	32.7	34.1	33.5
Share with BA degree or more	0.229	0.220	0.260	0.240
Non-White share	0.257	0.203	0.274	0.277
Foreign-born share	0.081	0.048	0.068	0.081

Notes: Industry employment shares, population, employment-population ratio, and real earnings per capita are measured two years before the recession start year. The last three variables are measured as of the closest decennial census to the recession start year, except for the 2007–2009 recession, which is measured from the 2005–2009 ACS. Population growth is from 1969 to 1973 for the 1973-1975 recession and over the ten years before the recession start for the other episodes. We define an area as experiencing a more severe recession if its log employment change for a given recession is less than the median across the 358 CBSAs for that recession.

Source: Authors' calculations of data from BEAR, decennial censuses and ACS (via IPUMS and NHGIS), and SEER.

that experienced a more severe recession. Controlling for prerecession population growth eliminates these pre-trends. Since population growth is calculated over the decade before the recession, it is likely we eliminate secular trends (such as growing migration to certain metros in the South and West).

The results in panel A of Figure 3 indicate that local employment losses during recessions are extremely persistent. The recession severity variable s_i^r is mechanically correlated with a drop in log employment during the recession. There is no mechanical relationship for the postrecession coefficients, however, which show little to no recovery over the subsequent ten years. Moreover, the confidence intervals imply that we can reject a return to initial peak employment in every postrecession year. The graph also shows that the persistent decline in employment is not affected by whether we control for the severity of previous recessions, and there is no evidence of subsequent

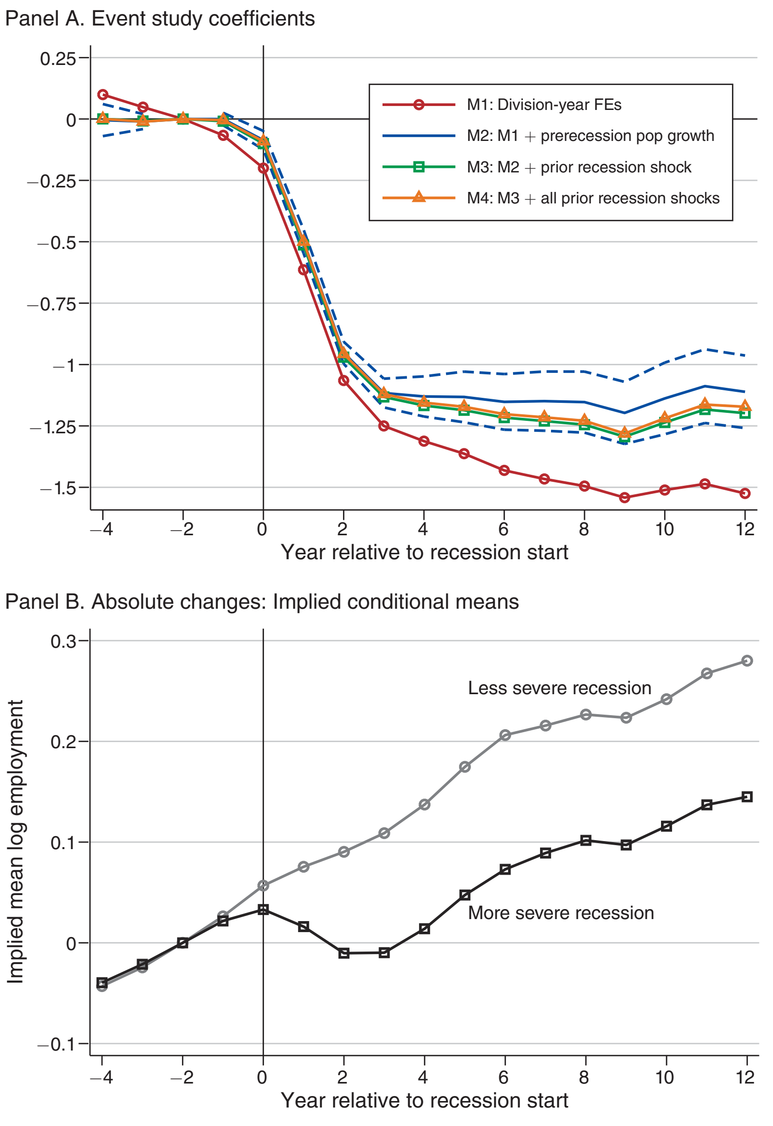

Figure 3. The Evolution of Metropolitan Area Log Employment after Recessions

Notes: Panel A reports estimates of equation (1). The dependent variable is log wage and salary employment from BEAR data, and the key independent variable is the change in log wage and salary employment during the recession from BEAR data. Specifications are indicated by the legend. There are 358 metropolitan areas in the sample. Standard errors are clustered by metropolitan area. In panel B, we use estimates of specification 2 to construct mean log employment for metropolitan areas with a more versus less severe recession (based on whether the log employment change is greater than or less than the median log employment change during the recession), holding all other covariates in the regression at their mean value. These conditional means are normalized to equal zero two years before the recession start. The average log employment change during the recession in the less severe recession group is 3.7 percent, and the average change in the more severe recession group is −4.9 percent.

Source: Authors' calculations using BEAR and SEER data.

discrete jumps, as might occur from a later shock. We obtain similar results when examining employment from County Business Patterns data, where we also see a persistent decline in the number of establishments (online Appendix Figure 4).

Panel B of Figure 3 illustrates how the relative changes identified by equation (1) translate into aggregate outcomes by displaying the implied evolution of mean log employment in metropolitan areas with a more versus less severe recession.[25] The postrecession level of employment is persistently lower in areas where the recession was more severe relative to areas where the recession was less severe. However, employment grows after the business cycle trough in absolute terms in both types of areas.

Panel A of Table 4 summarizes the (preferred) specification 2 results seven to nine years after the business cycle trough.[26] The employment elasticity is -1.1, which indicates that a 10 percent decrease in employment during the recession is followed by 11 percent lower employment 7–9 years later. Because recession severity varies both across recessions and across areas within a given recession (online Appendix Figure 2), we also report standardized coefficients. On average, a 1 standard deviation employment decline leads to a 6.6 percent decrease in employment 7–9 years after the trough.

The consequences of these decreases in employment depend on the degree of population response. We examine this next.

B. *Population*

Panel A of Figure 4 presents estimates of equation (1) where the dependent variable is the log of the total working age population (15+). For brevity, we show only the results from specification 2, although the patterns are robust to specifications 3 and 4. We find that areas with greater job loss experience postrecession decreases in population that double in magnitude over the postrecession period. The summary estimates in panel A of Table 4 indicate that a 10 percent decrease in employment during the recession is followed by a 5.8 percent decrease in population 7–9 years after the trough. For a 1 standard deviation employment decrease, this amounts to a 3.3 percent relative decrease in population.[27]

C. *Employment-Population Ratio*

Population declines by less than employment in areas that experience more severe recessions. This implies that the employment-population ratio falls after recessions in these areas. To examine this pattern more directly, we use the log of the ratio of employment to working age population as the dependent variable in equation (1).[28]

[25] We construct these conditional means using estimates of equation (1), holding all covariates besides recession severity at their mean value and defining the gap between a more and less severe recession as a log employment change difference of -0.086 (equal to the difference in mean recession severity for areas with a log employment change below or above the median).

[26] We generate the results in this table by pooling the coefficients in equation (1) for posttrough years seven through nine. Estimating a pooled coefficient summarizes the medium-term changes while also increasing precision. Online Appendix Table 6 presents results separately for each recession.

[27] Recession-specific estimates are in online Appendix Figure 5. Consistent with the previously documented decline in migration (Molloy, Smith, and Wozniak 2014; Dao, Furceri, and Loungani 2017), postrecession declines in population have become smaller over time.

[28] Our employment-population measure is the ratio of the count of jobs to the number of working age people; because of multiple job holdings, it is not strictly comparable to official employment-population ratios, which represent the share of the population that is employed.

Table 4—Summary of Changes in Metropolitan Area Economic Activity, Seven to Nine Years after Business Cycle Trough

Dependent variable	Coefficient on log employment decrease (1)	Implied change from 1 SD decrease in log employment (2)
Panel A: Dependent variables from BEAR and SEER		
Log employment	−1.141	−0.066
	(0.072)	
Log population age 15+	−0.577	−0.033
	(0.049)	
Log employment-population ratio	−0.564	−0.033
	(0.056)	
Log earnings per capita	−0.893	−0.052
	(0.078)	
Log earnings per worker	−0.329	−0.019
	(0.039)	
Panel B: log annual earnings, without composition adjustment		
Average log earnings	−0.394	−0.023
	(0.055)	
10th percentile, log earnings	−0.637	−0.037
	(0.105)	
50th percentile, log earnings	−0.350	−0.020
	(0.053)	
90th percentile, log earnings	−0.219	−0.013
	(0.040)	
Panel C: Weekly and hourly earnings		
Average log weekly earnings	−0.347	−0.020
	(0.046)	
Average log hourly earnings	−0.307	−0.018
	(0.042)	
Panel D: log annual earnings, with composition adjustment		
Average log earnings	−0.338	−0.020
	(0.048)	
10th percentile, log earnings	−0.518	−0.030
	(0.101)	
50th percentile, log earnings	−0.301	−0.017
	(0.041)	
90th percentile, log earnings	−0.261	−0.015
	(0.035)	

Notes: Table reports estimates of equation (1). Column 1 reports the point estimate and standard error, and column 2 contains the point estimate multiplied by the standard deviation of the log employment change during a recession (0.058). The dependent variable is indicated in the row. In panel A, the dependent variable is constructed as the change relative to two years before the nationwide business cycle peak, and we report the pooled coefficient for years seven to nine after the business cycle trough. In panels B to D, the dependent variable is constructed as the change between prerecession and postrecession years (1970 to 1980, 1980 to 1990, 1990 to 2000, 2000 to 2005–2007, and 2005–2007 to 2015–2017). The underlying sample for panels B to D is limited to individuals age 25–54 and then collapsed to 358 metropolitan areas. The dependent variables in panel D are constructed using residuals from regressing log earnings on indicators for education, age, sex, and race/ethnicity (White/Black/Hispanic/other) plus interactions between the education indicators and a quartic in age. The key independent variable is the change in log wage and salary employment during the recession from BEAR data. All regressions control for division-year fixed effects and interactions between prerecession population growth and year indicators. There are 358 metropolitan areas in the sample. Standard errors are clustered by metropolitan area.

Source: Authors' calculations using BEAR, SEER, decennial census, and ACS data.

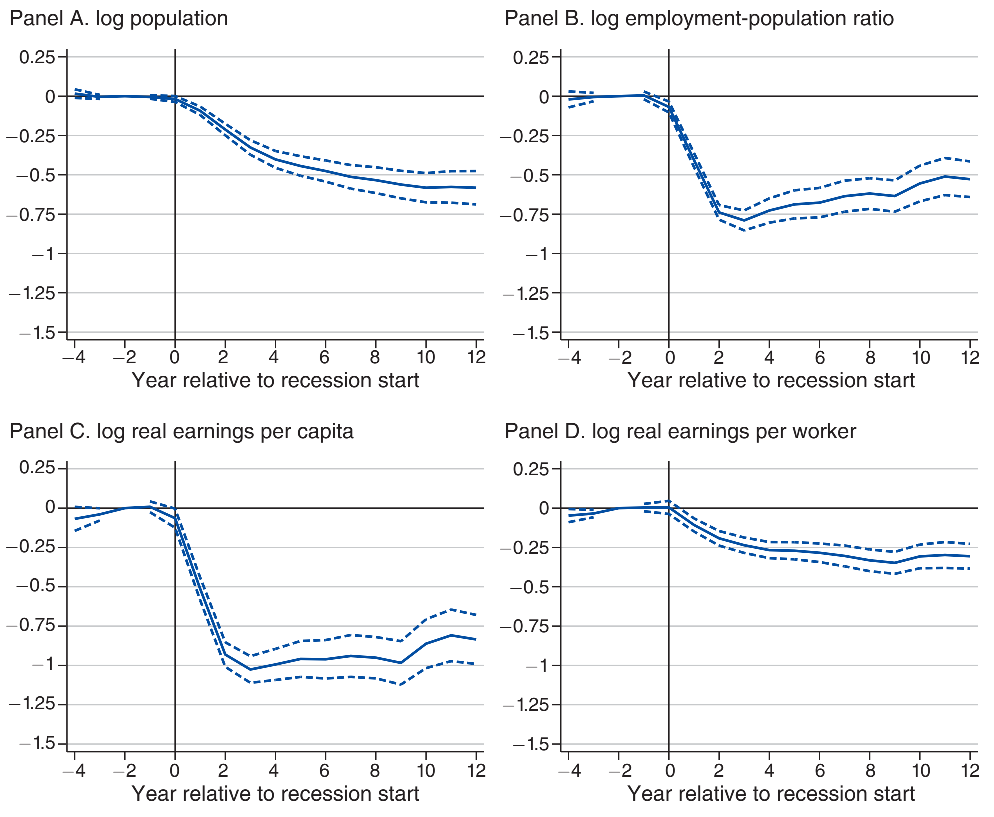

Figure 4. The Evolution of Metropolitan Area Log Population, Log Employment-Population Ratio, Log Real Earnings Per Capita, and Log Real Earnings per Worker after Recessions

Notes: Figure reports estimates of equation (1) for specification 2. The dependent variable is log population age 15 and above in panel A, the log ratio of wage and salary employment to population age 15 and above in panel B, log real earnings per capita (age 15+) in panel C, and log real earnings per worker in panel D. See notes to Figure 3.

Source: Authors' calculations using BEAR and SEER data.

Panel B of Figure 4 shows that the employment-population ratio falls during recessions and remains below the prerecession peak even a decade after a recession's end.[29] Due to the relatively flat employment trajectory and steady population decline, the employment-population ratio shows a slight recovery over time. As reported in panel A of Table 4, the average elasticity seven to nine years posttrough is about −0.6. Given a mean employment-population ratio of about 60 percent, this elasticity implies that a 10 percent decrease in employment during a recession is followed by a 3.4 percentage point decline in the employment-population ratio. A 1 standard deviation employment decline leads to a 3.3 percent (2.0 percentage point) decrease in the employment-population ratio.

[29] Recession-specific estimates are in online Appendix Figure 6.

The estimates in Table 4 facilitate a simple decomposition of the postrecession decline in employment, namely that the postrecession change in log employment equals the change in log population plus the change in the log employment-population ratio. On average, the decline in the employment-population ratio accounts for about half of the decline in employment seven to nine years after the business cycle trough, with the remaining half explained by the decline in population.

D. *Earnings per Capita*

Local employment losses could be followed by broader changes than a persistent decline in the employment-population ratio. For example, as explained in Section I, local labor markets could also face declines in wages, and reductions in employment could extend beyond the extensive margin to also affect hours worked. To understand the broader consequences of local employment losses, we examine changes in log earnings per capita. The results in panel C of Figure 4 show evidence of persistent reductions in earnings per capita following recessions.[30] The medium-term elasticity in panel A of Table 4 is -0.9, which implies that a 1 standard deviation greater employment decline is followed by a 5.2 percent larger relative decrease in earnings per capita 7–9 years after the trough.[31,32]

E. *Earnings per Worker*

Any reduction in wages and hours following local employment losses during recessions can also be examined through effects on log annual earnings per worker, which encapsulates both the quantity and quality of employment. Panel D of Figure 4 shows evidence of a persistent decline in earnings per worker that is sizable but smaller than the decrease in earnings per capita.[33] The definitions of the outcomes in panel A of Table 4 facilitate a decomposition of the decline in earnings per capita. In particular, the change in log earnings per capita equals the sum of the change in the log employment-population ratio and the change in log earnings per worker. We find that 63 percent of the postrecession decrease in earnings per capita is explained by the decline in the employment-population ratio, with the remaining 37 percent explained by the decrease in earnings per worker.

[30] Recession-specific estimates are in online Appendix Figure 7.

[31] Our preferred earnings measure includes wages, salaries, and supplements (benefits), which are available only by place of work. We show in online Appendix Figure 8 that our findings are not sensitive to the place of residence versus place of work distinction.

[32] Recessions could also lower housing prices. Using housing price indices from the Federal Housing Finance Agency (Federal Housing Finance Agency 1975–2019), we find evidence of a relative decrease in housing prices in areas that experience larger employment losses during a recession (online Appendix Figure 9, panel A). To explore whether these changes in prices offset the decline in earnings per capita, we combine these estimates with the two approaches of constructing local CPI from Moretti (2013). We find that 70–80 percent of the change in log earnings per capita remains after adjusting for local prices, as shown in panel B of online Appendix Figure 9.

[33] The numerator of this outcome is the same as in earnings per capita, but we change the denominator to the annual employment count rather than the working age population. Recession-specific estimates are in online Appendix Figure 10.

F. *Robustness*

Our results are robust to modifying the empirical specification in several different ways. In particular, online Appendix B.1 shows that our results are very similar when using private wage and salary employment from BEAR or QCEW data to measure recession severity. Online Appendix B.2 discusses results that isolate variation in the log employment change during recessions that is predicted by an area's preexisting industrial specialization (Bartik 1991). These estimates reveal persistent declines in local economic activity that are similar in magnitude to our main results, which implies that our finding of a persistent postrecession decline in local economic activity is not driven entirely by idiosyncratic shifts in local labor demand. The similarity of these results also suggests that the employment change during recessions is not driven by immediate endogeneous policy responses. Indeed, the event study results are also robust to replacing division-year fixed effects with state-year fixed effects (online Appendix Figure 14), which further address potential endogenous policy responses. Finally, online Appendix B.3 shows that our results are nearly identical when examining commuting zones instead of metropolitan areas.

IV. Discussion

Our results point to a different understanding of local labor market dynamics compared to papers that estimate rapid recovery in response to labor demand shocks (e.g., Blanchard and Katz 1992; Dao, Furceri, and Loungani 2017). While our finding of persistent declines in employment and population are qualitatively similar to these papers' findings, we provide evidence of persistent declines in the employment-population ratio, earnings per capita, and earnings per worker. This implies that labor demand shocks have much longer lasting consequences for local areas and that migration plays a smaller role in equilibrating local labor market outcomes. The literature studying the impacts of Bartik shocks over longer horizons also supports this view of the labor market (Bartik 1991; Bound and Holzer 2000; Notowidigdo 2020).

In this section, we present additional evidence that supports this interpretation and provides additional context. We demonstrate that our results are not driven by secular changes associated with areas' prerecession industrial specialization or demographic or labor market characteristics. We also show that all sectors experience a relative decline in employment, while population falls primarily because of lower in-migration. Moreover, the decrease in earnings among individuals who remain employed is explained mainly by a reduction in hourly wages as opposed to hours of work. Finally, we show that persistent local labor market declines do not simply reflect changes in the composition of residents.

A. *Supporting Evidence*

A possible concern is that our estimates simply reflect the effects of secular changes in the economy, such as the decline in manufacturing. This issue is closely related to the hypothesis of Amior and Manning (2018), who argue that slow regional recoveries are partly due to serially correlated labor demand shocks, which could resemble secular changes in annual data. Several factors point against these interpretations in our setting.

Most importantly, there is little evidence that the persistent decline in local economic activity is driven by subsequent shocks that occur after recessions. If areas faced a severe recession and then a serially correlated shock a few years later, we would expect to see postrecession years with sharp decreases in employment. These sharp changes are not evident in Figure 3. Instead, employment declines rapidly during the recession and then remains relatively flat over the following decade.[34] These results suggest that serially correlated labor demand shocks play a minor role in our setting.

To explore this issue further, we estimate regressions that additionally control for interactions between recession-specific year indicators and prerecession metro area characteristics. One set of regressions controls for shares of employment in each of ten sectors: agriculture, construction, finance, government, manufacturing, mining, retail trade, services, utilities, and wholesale trade. These controls absorb changes in economic activity that are associated with industrial specialization. For example, areas that specialize in manufacturing might have experienced reductions in employment for the past 50 years, due either to secular change or repeated shocks. Another set of regressions controls for the prerecession labor market and demographic characteristics examined in Table 3 and online Appendix Table 2, which could also correlate with preexisting trends or future nonrecession shocks. The results from both of these specifications, shown in online Appendix Figure 16, are similar to our baseline results from Figures 3 and 4. Our estimates of persistent postrecession declines do not simply reflect secular changes determined by industry structure or other labor market characteristics.

B. *Contextualizing Evidence*

Employment Declines across All Sectors.— Are the employment losses shown in Figure 3 broad-based or concentrated in certain industries? To explore this question, Figure 5 shows estimates of equation (1), where the dependent variable is log employment in each sector. For simplicity and ease of presentation, we present estimates for specification 2 only and suppress confidence intervals.[35] We find that the relative decline in employment is pervasive across sectors. Construction and manufacturing experience the largest short-term decreases, while government employment generally falls the least. The remaining industries tend to move similarly and lie in between; with the exception of construction, there is little evidence of an upward slope to suggest an eventual recovery in employment.[36]

Population Declines through Lower In-Migration.— What explains the decline in population? We use the SOI data to examine this question for the two most recent recessions. In particular, we decompose the net change in population into changes

[34] As shown in online Appendix Figure 3, this pattern generally holds for individual recessions as well.

[35] Online Appendix Figure 17 shows recession-specific estimates.

[36] We exclude agriculture and mining, which are small (especially in metropolitan areas) and highly spatially concentrated.

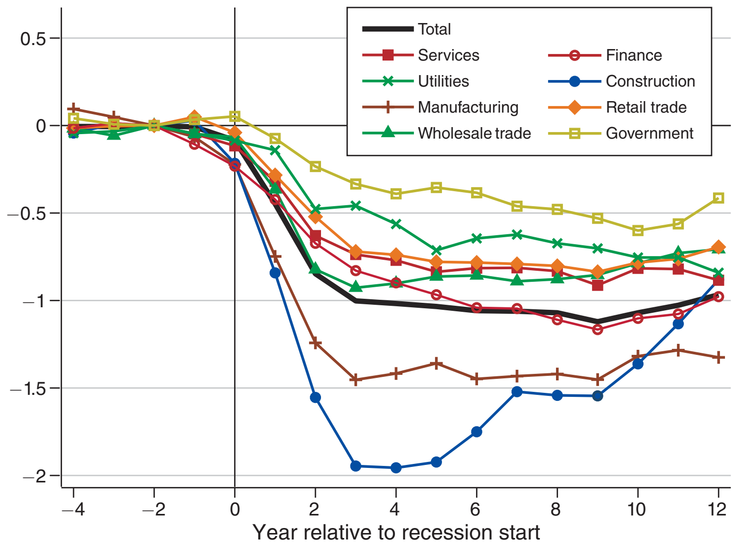

Figure 5. The Evolution of Metropolitan Area Log Employment by Sector after Recessions

Notes: Figure reports estimates of equation (1) for specification 2. The dependent variable is log employment from the indicated sector. We use BEAR data for the 1973–1975, 1980–1982, 1990–1991, and 2007–2009 recessions. We use QCEW data for the 2001 recession (due to SIC-NAICS industry seaming issues), except for government, which comes from BEAR. See notes to Figure 3.

Source: Authors' calculations using BEAR, SEER, and QCEW data.

in in-migration, out-migration, and residual net births. This decomposition starts with the identity

$$pop_{i,t} = pop_{i,t-1} + inmig_{i,t} - outmig_{i,t} + netbirths_{i,t}, \tag{2}$$

where $pop_{i,t}$ is population in location i and year t, $inmig_{i,t}$ is the number of in-migrants between year $t-1$ and t, $outmig_{i,t}$ is the number of out-migrants, and $netbirths_{i,t}$ is the number of births minus deaths. Iterating equation (2) and normalizing by a baseline population level two years before the recession start, we can decompose the proportional change in population from year $p(r)-2$ to year t into components for in-migration, out-migration, and net births as follows:

$$\frac{pop_{i,t}}{pop_{i,p(r)-2}} - 1 = \sum_{j=p(r)-1}^{t} \frac{inmig_{i,j}}{pop_{i,p(r)-2}} - \sum_{j=p(r)-1}^{t} \frac{outmig_{i,j}}{pop_{i,p(r)-2}} + \sum_{j=p(r)-1}^{t} \frac{netbirths_{i,j}}{pop_{i,p(r)-2}}. \tag{3}$$

As a starting point, panel A of Figure 6 shows that using the number of personal exemptions from the SOI data to construct the variable on the left-hand side of equation (3) yields results that are similar to the change in log population shown

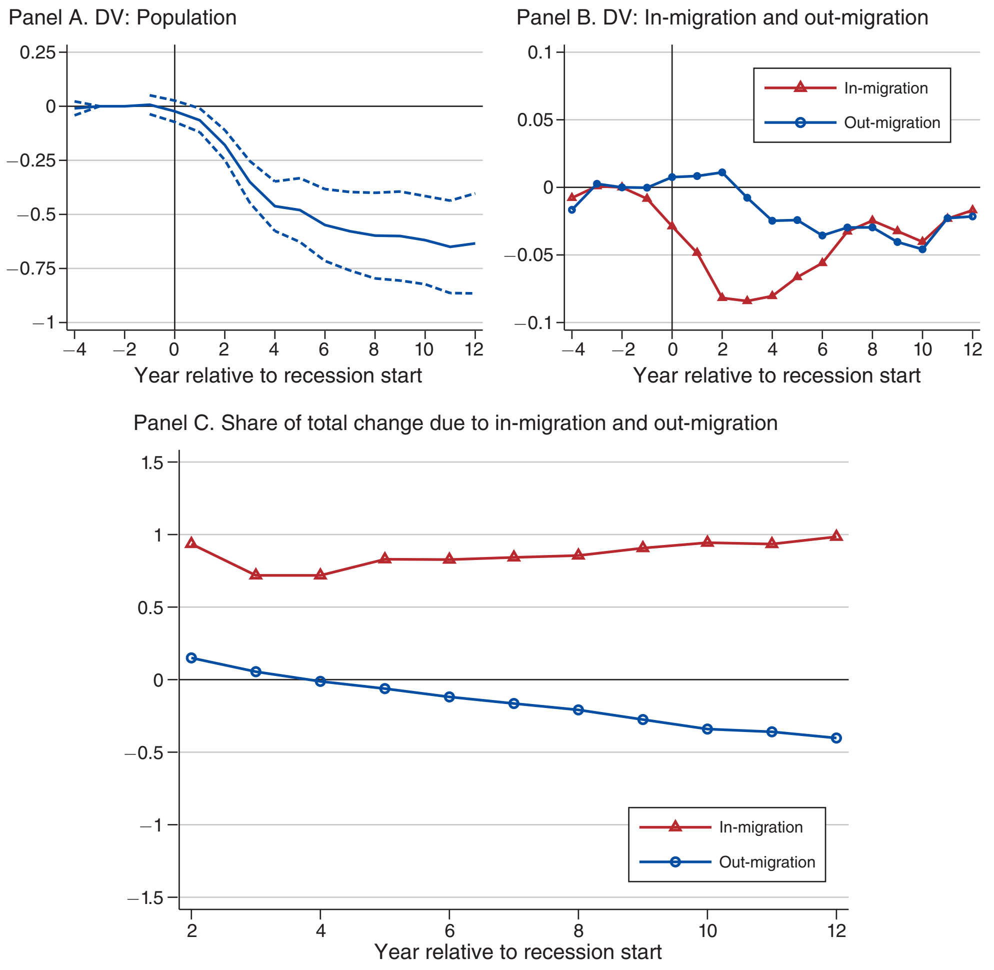

Figure 6. The Evolution of Metropolitan Area In-Migration and Out-Migration After Recessions

Notes: Figure reports results that stack the 2001 and 2007–2009 recessions and estimate a variant of equation (1) in which the dependent variable is the outcome in year t and we control for interactions between recession-specific year fixed effects and in-migration, out-migration, and net birth rates in year $p(r) - 2$. This approach facilitates an exact decomposition using the regression coefficients (including net births, which we do not show for brevity). In panel A, the dependent variable is the number of exemptions in year t divided by the same variable in year $p(r) - 2$. In panel B, the dependent variables are in-migration and out-migration relative to the number of exemptions in year $p(r) - 2$. In panel C, we divide cumulative sums of the coefficients from panel B by the coefficients in panel A; we multiply the out-migration coefficient by -1 so that a positive number indicates that a given population component contributes to the postrecession population decline. Regressions also include specification 2 controls described in the notes to Figure 3.

Source: Authors' calculations using CBP, BEAR, and SOI data.

in panel A of Figure 4.[37] Panel B of Figure 6 presents results where the dependent variables are annual migration inflows and outflows divided by the total number

[37] For this analysis, we stack the 2001 and 2007–2009 recessions and estimate a variant of equation (1) in which the dependent variable is an outcome as of year t, and we control for interactions between recession-specific year fixed effects and in-migration, out-migration, and net birth rates in year $p(r) - 2$. This approach facilitates an exact decomposition using the regression coefficients, although we omit the effect on net birth rates from the figures for brevity.

of exemptions in year $p(r) - 2$. By business cycle trough, in-migration rates have fallen sharply, with a 10 percent decrease in employment during the recession being followed by a reduction in annual in-migration of about 0.8 percent of prerecession population. Over the subsequent decade, in-migration rates gradually recover but remain depressed ten years after the business cycle trough. Out-migration shows little response until after the recession has ended. Beginning in the year after the trough, however, out-migration rates steadily decline for several years, with similar medium-term magnitudes as for in-migration.

To understand how these components contribute to the change in population, we use the decomposition in equation (3). In particular, we construct cumulative sums of the coefficients in panel B and divide these sums by the respective estimates in panel A. When we also multiply the out-migration estimates by -1, the three transformed coefficients—in-migration, out-migration, and net births—sum to 1 and fully decompose the postrecession population change in each period. The results in panel C reveal that lower in-migration accounts for essentially all of the medium-run decrease in population after recessions.[38] In contrast to a story of individuals moving away from places where recessions are more severe, the decrease in out-migration dampens the population decline.[39] The lack of out-migration is a natural explanation for why the population response is incomplete.

Earnings Decline throughout the Distribution, *via Lower Hourly Wages*.— We use census/ACS data to examine changes in the distribution of prime-age workers' earnings. Specifically, we estimate a variant of equation (1) in which the dependent variable is a pre-postrecession change.[40] We examine the mean and the tenth, fiftieth, and ninetieth percentiles of the log annual earnings distribution. The first row of panel B of Table 4 shows that estimates for mean log earnings are similar to those from the BEAR data on log earnings per worker. The percentile estimates in the next three rows indicate that earnings fall throughout the distribution, with larger changes at lower percentiles. These results are consistent with the finding that lower-earning demographic groups are more affected during recessions (Hoynes, Miller, and Schaller 2012).

Does the reduction in earnings stem from a reduction in hours worked, a reduction in earnings per hour, or both? To answer this question, we use the census/ACS data to estimate regressions where the dependent variable is the change in average log annual, weekly, or hourly earnings. If the earnings losses are driven by a reduction in hours, hourly wages could be relatively unaffected several years later. On the other hand, if the recession slows wage growth or displaced workers are less likely

[38] There is a decrease in net births (not shown) that offsets the decline in out-migration in explaining the net population decline.

[39] Monras (2020) also finds this pattern of relative population decline due to falling in-migration for the Great Recession, using variation in recession severity based on prerecession per capita debt and the share of employment in nontradable industries (see also Mian, Rao, and Sufi 2013). His calibrated general equilibrium model predicts that migration dissipates about 60 percent of the long-term impact on wages following the Great Recession.

[40] We use the 1970 and 1980 censuses for the 1973–1975 recession, the 1980 and 1990 censuses for the 1980–1982 recession, the 1990 and 2000 censuses for the 1990–1991 recession, the 2000 census and 2005–2007 ACS for the 2001 recession, and the 2005–2007 and 2015–2017 ACS for the 2007–2009 recession. Because the variables used are based on the previous calendar year (census) or preceding 12 months (ACS), these changes straddle the periods when recessions occur.

to find good employer matches, hourly wage losses may explain more of the annual earnings declines. The results in panel C of Table 4 indicate that the latter story better fits the data and accord with Lachowska, Mas, and Woodbury (2020), as the estimated decline in log hourly wages explains about three-quarters of the decline in log annual earnings. Decreases in work attachment at the intensive margin therefore explain relatively little of the persistent reduction of annual earnings among individuals who remain employed.[41,42]

The Role of Changes in the Composition of Residents.— A remaining explanation for why recessions are followed by persistent declines in the employment-population ratio and earnings per capita is a change in worker composition due to differential migration responses. For example, if highly educated workers are more likely to leave an area in response to a decline in employment (Bound and Holzer 2000; Wozniak 2010; Notowidigdo 2020), then average wages might fall because of a change in worker composition. Composition shifts are not a threat to our identification strategy because our unit of analysis is an area rather than an individual, but they are an interesting mechanism to understand.

To quantify the role of composition shifts, we examine changes in residualized earnings. We regress log annual earnings of prime-age workers from the census and ACS on indicators for education (of which there are 11), age (30), sex (2), and race/ethnicity (4), plus interactions between the education indicators and a quartic in age. We estimate these regressions separately for each year and use metro area averages and percentiles of the residuals as dependent variables in our regressions. Panel D of Table 4 presents results for composition-adjusted wage and salary earnings. The composition-adjusted results tend to be somewhat smaller in magnitude, which indicates that composition shifts partly contribute to the persistent decline in earnings. However, the composition-adjusted estimates are still at least 80 percent as large as the unadjusted ones. This finding suggests that the persistent postrecession declines in average earnings are not primarily driven by changes in worker characteristics correlated with these variables.

The availability of annual population estimates by age from the SEER data allows us to use a complementary approach to explore the role of shifts in the age distribution of residents in each postrecession year. In particular, we predict the average change in the log employment-population ratio due to changes in the age structure by combining estimates of the postrecession evolution of the share of the population age 0–14, 15–24, 25–34, 35–44, 45–54, 55-64, and over 65 with the cross-sectional, prerecession relationship between the age structure and the log employment-population ratio.[43] The results in panel A of online Appendix Figure 19 show that changes in the age

[41] These results do not conflict with our finding that the reduction in the employment-population ratio explains most of the decline in earnings per capita because our analysis of census/ACS data conditions on earnings being positive.

[42] A potential concern is that the census/ACS results in panels B and C of Table 4 are difficult to compare to the results using BEA data in panel A because of differences in when outcomes are measured. This issue is of limited importance in practice, as results that use the BEA data for the census/ACS years are very similar to the baseline BEA results, as indicated in online Appendix Table 8.

[43] We cannot use the same set of observed variables in this annual approach as we use with the long difference for the census/ACS data because the annual SEER data lack population counts by education.

structure predict a decrease in the log employment-population ratio that is equal to 40 percent of the actual long-run decrease. The results for log earnings per capita, which are constructed analogously and shown in panel B, are similar. Panel C shows that these results arise from a decrease in the share of population below age 45.[44] In line with the results using individual-level data on prime-age workers' earnings from the census and ACS, these findings suggest that shifts in the composition of residents explain some but not all of the persistent decline in local labor market outcomes after recessions.

Long-Run Results.— Our main results follow local labor markets for a dozen years after recession start. Do local areas eventually recover over a longer horizon? Online Appendix Figures 20 and 21 show that neither employment nor employment-population ratios had recovered by 2019 for *any* recession.

V. A Comparison to Results from the Blanchard and Katz (1992) Model

Our finding that recessions are followed by persistent declines in the employment-population ratio and earnings per capita differs from the well-known results of BK, which imply that the unemployment rate, the labor force participation rate, the employment-population ratio, and wages return to trend within ten years after state-level employment declines. Our empirical strategy is fundamentally similar to BK in that we both rely on cross-sectional variation in how local areas respond to employment changes. The key difference is that BK, and the many papers that follow their approach, estimate VARs and then calculate impulse response functions (IRFs), while we estimate regression models that impose no constraints on how coefficients vary across years relative to the recession start. This section explores why our results differ.

To facilitate discussion, we first introduce the BK VAR. The key variables are the annual change in log employment, $\Delta e_{i,t}$, the level of the log employment–labor force ratio, $el_{i,t}$, and the level of the log labor force–working age population ratio, $lp_{i,t}$. BK account for nationwide trends by differencing out the same variables for the aggregate US economy. They estimate the following recursive VAR using state-level data from 1976–1990:

$$\Delta e_{i,t} = \alpha_{i10} + \alpha_{11}(L)\Delta e_{i,t-1} + \alpha_{12}(L)el_{i,t-1} + \alpha_{13}(L)lp_{i,t-1} + \epsilon_{i,e,t}, \tag{4}$$

$$el_{i,t} = \alpha_{i20} + \alpha_{21}(L)\Delta e_{i,t} + \alpha_{22}(L)el_{i,t-1} + \alpha_{23}(L)lp_{i,t-1} + \epsilon_{i,el,t}, \tag{5}$$

$$lp_{i,t} = \alpha_{i30} + \alpha_{31}(L)\Delta e_{i,t} + \alpha_{32}(L)el_{i,t-1} + \alpha_{33}(L)lp_{i,t-1} + \epsilon_{i,lp,t}. \tag{6}$$

BK include two lags of each explanatory variable along with state fixed effects α_{i10}, α_{i20} and α_{i30}. After estimating these equations (which can be done using three

[44] Our results are similar to Cajner, Coglianese, and Montes (2021) in documenting a decrease in the share of population between ages 25 and 44 and an increase in the share of population at older ages. There also are some differences: we find a decrease in the share of population ages 15–24, while their results suggest a slight increase. Given the differences in the unit of analysis, methodology, and source of identifying variation, we see the results on the relative responses of different age groups as being quite consistent with each other.

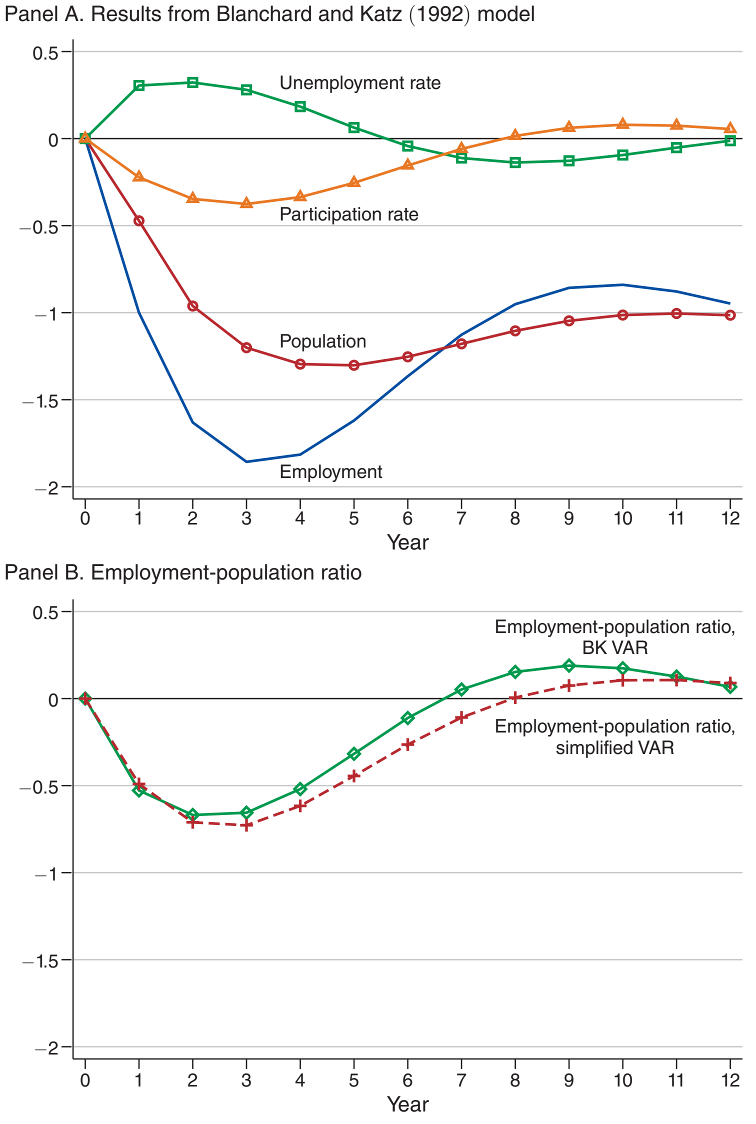

Figure 7. Impulse Response Functions to Negative Log Employment Shock from Vector Autoregressions

Notes: Figure shows IRFs of indicated variables with respect to a negative log employment shock. We construct IRFs for the BK VAR using estimates of equations (4)–(6). For the simplified VAR in panel B, we use equations (7)–(8). Sample contains 50 states and Washington, DC, from 1976–1990.

Source: Authors' calculations using BLS CES and LAUS data.

separate OLS regressions), BK construct the IRFs of each variable with respect to a 1 percent decrease in employment (i.e., a reduction in $\epsilon_{i,e,t}$ of 0.01).[45] Primary interest lies in these IRFs, which are constructed using only the coefficients in equations (4)–(6).

Figure 7 shows IRFs of log employment, the "unemployment rate" (one minus the log employment–labor force ratio), the log participation rate, and log population.

[45] Because this is a recursive VAR, there is a natural unit of measurement for $\epsilon_{i,e,t}$. In contrast, a structural VAR does not feature this property (see, e.g., Stock and Watson 2001).

We use BLS data from 1976–1990 to generate these results, which are extremely similar to Figure 7 of BK.[46] Notably, the unemployment rate and participation rate completely recover within eight years.

Our preferred unit of geography is a metropolitan area. When using substate areas, reliable data on labor force participation are available for a limited time period at best.[47] Consequently, the most comparable outcome is the employment-population ratio. The IRF of the log employment-population ratio can be constructed as the sum of the IRFs of the log employment–labor force ratio and the log labor force–population ratio. Panel B of Figure 7 shows this IRF from the BK model. As expected given the results in panel A, the IRF shows complete recovery of the employment-population ratio.

To facilitate the analysis below, we simplify the BK model in two ways. First, we estimate a two-equation VAR in first-differences of log employment and levels of the log employment-population ratio, $ep_{i,t}$. Second, we include only one lag of each variable. The resulting VAR is

$$\Delta e_{i,t} = \tilde{\alpha}_{i10} + \tilde{\alpha}_{11} \Delta e_{i,t-1} + \tilde{\alpha}_{12} ep_{i,t-1} + \tilde{\epsilon}_{i,e,t}, \tag{7}$$

$$ep_{i,t} = \tilde{\alpha}_{i20} + \tilde{\alpha}_{21} \Delta e_{i,t} + \tilde{\alpha}_{22} ep_{i,t-1} + \tilde{\epsilon}_{i,ep,t}. \tag{8}$$

These simplifying assumptions have little impact on the estimated IRF of the log employment-population ratio, as shown in panel B of Figure 7.

Equations (7) and (8) permit simpler expressions of the IRF in terms of the underlying parameters. Consider a one-time change in log employment in period t through $\tilde{\epsilon}_{i,e,t}$. The subsequent impacts on the log employment-population ratio in periods t through $t + 2$ are

$$\frac{dep_{i,t}}{d\tilde{\epsilon}_{i,e,t}} = \tilde{\alpha}_{21}, \tag{9}$$

$$\frac{dep_{i,t+1}}{d\tilde{\epsilon}_{i,e,t}} = \tilde{\alpha}_{21}^2 \tilde{\alpha}_{12} + \tilde{\alpha}_{21} \tilde{\alpha}_{11} + \tilde{\alpha}_{21} \tilde{\alpha}_{22}, \tag{10}$$

$$\begin{aligned} \frac{dep_{i,t+2}}{d\tilde{\epsilon}_{i,e,t}} = {} & \tilde{\alpha}_{21}^3 \tilde{\alpha}_{12}^2 + 2\tilde{\alpha}_{21}^2 \tilde{\alpha}_{11} \tilde{\alpha}_{12} + 2\tilde{\alpha}_{21}^2 \tilde{\alpha}_{22} \tilde{\alpha}_{12} + \tilde{\alpha}_{21} \tilde{\alpha}_{11}^2 \\ & + \tilde{\alpha}_{21} \tilde{\alpha}_{22}^2 + \tilde{\alpha}_{21} \tilde{\alpha}_{11} \tilde{\alpha}_{22}. \end{aligned} \tag{11}$$

Similar expressions exist for the IRF at later horizons, but these first few periods are adequate to highlight some important takeaways. First, bias in the OLS estimates of equations (7) and (8) can generate bias in the IRF because the IRF is a function

[46] We follow BK in measuring employment using the BLS Current Employment Statistics. We also follow the same approach as BK in using the BLS Local Area Unemployment Statistics (LAUS) (Bureau of Labor Statistics 1976–2019) to measure the number of individuals that are unemployed or not in the labor force and then constructing population as the sum of employment, unemployment, and not-in-labor-force counts.

[47] The BLS provides county-level labor force estimates from 1990 onward. A separate series contains county-level labor force estimates from 1976–1989, but the BLS stresses that this series is "unofficial" and not comparable to the 1990-and-forward series. Both datasets rely substantially on extrapolations from statistical models, as household surveys are not large enough to reliably measure unemployment and labor force for most counties.

of the coefficients in these equations. Second, bias in the IRF can increase in importance over time. For example, if the OLS estimates are attenuated, this bias generates an IRF that can converge toward zero even if the true IRF does not. This arises because the exponents in the IRF and the interactions between parameters increase with time, potentially magnifying bias.[48]

The potential for finite sample bias in autoregressive models, including VARs, has long been recognized (e.g., Hurwicz 1950; Shaman and Stine 1988; Stine and Shaman 1989; Pope 1990; Hall 1992; Kilian 1998, 1999; Kilian and Lütkepohl 2017).[49] This bias arises because residuals are not independent of all regressors in an autoregression since regressors are lagged dependent variables.[50]

To explore this issue further, we conduct a Monte Carlo study of finite sample bias. We focus on a data-generating process (DGP) where a decrease in employment leads to a persistent reduction in the employment-population ratio. We do not argue that this is the true DGP. Instead, this exercise illustrates how the BK VAR can fail to estimate a persistent decline in the employment-population ratio when one is actually present. For the Monte Carlo exercise we assume that log employment is a random walk,

$$e_{i,t} = e_{i,t-1} + \varepsilon_{i,e,t}, \tag{12}$$

and that log population depends on changes in log employment as follows:

$$p_{i,t} = p_{i,t-1} + \left(1 - \phi\right)\Delta e_{i,t} + \varepsilon_{i,p,t}. \tag{13}$$

This implies that the log employment-population ratio is

$$ep_{i,t} = ep_{i,t-1} + \phi\Delta e_{i,t} - \varepsilon_{i,p,t}. \tag{14}$$

In terms of equations (7) and (8), this DGP sets $\tilde{\alpha}_{i10} = \tilde{\alpha}_{i20} = 0$ (state fixed effects do not matter), $\tilde{\alpha}_{11} = \tilde{\alpha}_{12} = 0$ (log employment is a random walk), $\tilde{\alpha}_{21} = \phi$, and $\tilde{\alpha}_{22} = 1$. Changes in log employment have a permanent effect on the log employment-population ratio, with the true IRF equal to ϕ at all horizons.[51]

We calibrate the DGP using state-level LAUS data. We assume that all variables are distributed normally. The first period mean and variance of $e_{i,t}$ and $p_{i,t}$ equal those observed in the 1976 LAUS data, and the variances of $\varepsilon_{i,e,t}$ and $\varepsilon_{i,p,t}$

[48] More generally, if $a \in \left(0,1\right)$ is an attenuation factor, then $\left(ax\right)^t$ converges to zero faster than x^t.

[49] Kilian (1998, 1999) specifically addresses bias in impulse responses. The methods discussed in these papers allow for bias-corrected confidence intervals of impulse responses, but we focus on point estimates here for simplicity. In general, "there is no consensus in the literature that impulse responses should be estimated based on bias-adjusted slope parameters rather than the original [least squares] estimates" (Kilian and Lütkepohl 2017, p. 37).

[50] In his discussion of BK, Hall (1992) raises a concern about finite sample bias but speculates that such bias does not drive BK's conclusions. Amior and Manning (2018) theorize that the limited number of lags in the BK model could explain why BK find faster recovery than Amior and Manning (2018). Bias caused by a limited number of time periods—which we explore here—is distinct from whether the VAR has the appropriate lag structure.

[51] The true IRF for employment is 1, and the true IRF for population is $1 - \phi$.

approximate the variance of log employment and population in subsequent years.[52] We focus on the case where $\phi = 0.75$, with 50 cross-sectional observations and different time-series lengths, T. We study the response to a decrease in $\varepsilon_{i,e,t}$ as in BK.

Panel A of Figure 8 plots the true IRF for the employment-population ratio along with average estimates of the IRF across 499 Monte Carlo simulations. The true IRF reveals a persistent decrease in the employment-population ratio following a one-time decrease in employment. For $T = 15$, which is approximately the number of years available to BK when they wrote their paper, finite sample bias leads to rapid recovery of the employment-population ratio. Ten years after the shock, the IRF estimate is downward biased (in absolute value) by 89 percent. This bias remains very large for $T = 25$ and $T = 50$. Because previous work on local labor markets uses annual data, the relevant value of T ranges from 15 to 50. For $T = 100$ the bias remains sizable, at 25 percent 1 decade after the shock. Even for $T = 500$, finite sample bias incorrectly implies a gradual recovery.[53] The bias stems from an insufficient number of time series observations, so instrumental variables, which rely on asymptotic consistency, do not solve this problem in general. Indeed, we find that a sufficiently strong instrumental variable (as has been used in previous work) generates nearly identical results in our DGP (in which an instrument is not needed to obtain consistent estimates).

Finite sample bias also affects the estimated IRFs for other variables in the VAR, as shown in online Appendix Figure 25. In the simplified version of the BK VAR, finite sample bias incorrectly implies too small of a decline in employment and too large of a decline in population. Finite sample bias for the employment-population ratio is more severe than for employment because the IRF for the employment-population ratio depends on more parameters, each of which suffers from finite sample bias.[54] The opposite sign of the bias for population is a result of the structure of the VAR.[55]

Regressions that mirror our preferred specification in equation (1) do not suffer from finite sample bias due to small T in this setting. To show this, we use the same DGP and estimate the following regression:

$$ep_{i,t} - ep_{i,0} = -1 \times \Delta e_i \delta_t + \beta_t + \varepsilon_{i,t}, \tag{15}$$

[52] In particular, we set $e_{i,0} \sim \mathcal{N}(13.88, 1.03^2)$, $p_{i,0} \sim \mathcal{N}(14.43, 1.05^2)$, $\varepsilon_{i,e,t} \sim \mathcal{N}(0, 0.015^2)$, and $\varepsilon_{i,p,t} \sim \mathcal{N}(0, 0.015^2)$.

[53] Online Appendix Table 19 reports the underlying bias in estimates of the parameters of equations (7) and (8) for various values of T. All parameters are biased. While this bias is modest in many cases, it is amplified in the IRF. The IRF bias is of primary interest because the IRF is used to quantify the extent of recovery.

[54] For example, the IRFs of employment ($e_{i,t}$) and the employment-population ratio ($ep_{i,t}$) in the first period are $de_{i,t}/d\tilde{\epsilon}_{i,e,t} = 1$ and $dep_{i,t}/d\tilde{\epsilon}_{i,e,t} = \tilde{\alpha}_{21}$. There is no bias in the first-period IRF for employment, but there is bias for the employment-population ratio. In the second period, the IRFs for employment and the employment-population ratio are $de_{i,t+1}/d\tilde{\epsilon}_{i,e,t} = 1 + \tilde{\alpha}_{11} + \tilde{\alpha}_{12}\tilde{\alpha}_{21}$ and $dep_{i,t+1}/d\tilde{\epsilon}_{i,e,t} = \tilde{\alpha}_{21}^2\tilde{\alpha}_{12} + \tilde{\alpha}_{21}\tilde{\alpha}_{11} + \tilde{\alpha}_{21}\tilde{\alpha}_{22}$. This pattern holds generally and is the result of employment being the key "shock" variable in the VAR system.

[55] In the BK VAR, the IRF for population is inferred from the response of the other variables. In the simplified version of the BK model, the IRF for population equals the IRF for employment minus the IRF for the employment-population ratio. As a result, the greater bias in the IRF for the employment-population ratio implies that the population IRF is biased in the opposite direction of the employment IRF.

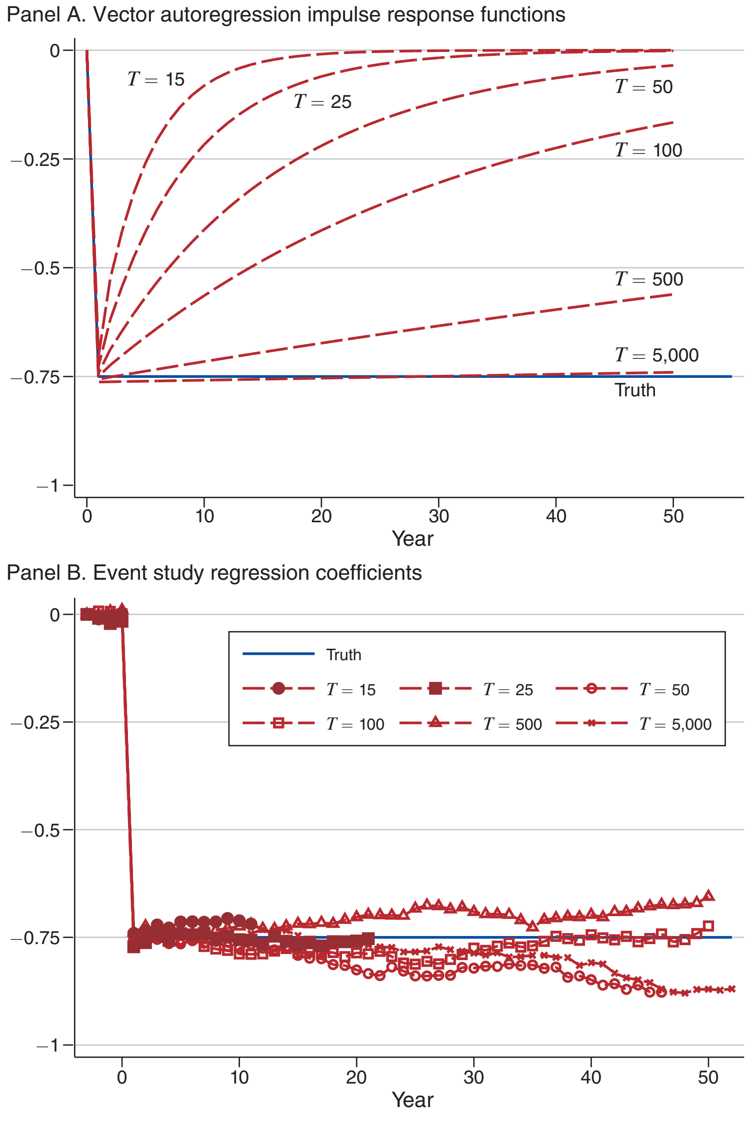

Figure 8. Comparison of Finite Sample Bias from Vector Autoregression Impulse Response Functions and Event Study Regressions for the Log Employment-Population Ratio

Notes: Panel A displays average estimates of IRFs of the log employment-population ratio with respect to a negative log employment shock based on estimates of equations (7)–(8). Panel B displays average estimates of δ_t from the regression in equation (15). For both panels, we simulate data following equations (12)–(14). We set $e_{i,0} \sim \mathcal{N}(13.88, 1.03^2)$, $p_{i,0} \sim \mathcal{N}(14.43, 1.05^2)$, $\varepsilon_{i,e,t} \sim \mathcal{N}(0, 0.015^2)$, $\varepsilon_{i,p,t} \sim \mathcal{N}(0, 0.015^2)$, $\phi = -0.75$, and $N = 50$. Results are based on 499 Monte Carlo simulations.

where the change in log employment Δe_i occurs between years 0 and 1, we multiply this change by -1 to mirror our analysis elsewhere, and β_t is a year fixed effect. To be consistent with the VAR IRFs, we normalize the coefficient $\delta_0 = 0$. This is the direct analog of equation (1). Under this DGP, we have $\delta_t = -0.75$ for all years $t \geq 1$. Hence, the true values of the coefficient δ_t and the IRF for the employment-population ratio coincide for all years after the measured log employment change. Panel B of Figure 8 shows that there is no systematic bias in estimates of δ_t regardless of T.[56]

In sum, finite sample bias can lead the BK VAR to find evidence of recovery when there is none. The regressions that we estimate are not subject to this finite sample bias in empirically relevant DGPs. We believe that this is the main explanation for why we find widespread evidence of persistent declines in employment-population ratios and earnings per capita while papers estimating the BK VAR generally do not.[57] To be clear, we do not claim that all VARs are incapable of identifying persistent changes. However, finite sample bias is evident in DGPs that are relevant for VARs estimated in previous work on local labor markets.

VI. Conclusion

Studying recessions over the course of 50 years, this paper shows that local employment losses that emerge during recessions are followed by long-lasting relative declines in employment, population, employment-population ratios, and earnings per capita. These patterns are consistent with harder-hit areas facing a persistent decline in labor demand relative to other areas, with labor supply being insufficiently responsive to restore prerecession employment-population ratios and wages. One explanation for why these results have not been shown before is that an influential approach in the literature—estimating VARs and calculating IRFs as in BK—can incorrectly find convergence after a persistent decline in local employment because of finite sample bias. In contrast, the regressions that we estimate do not suffer from this bias.

Cross-sectional variation in recession severity allows us to estimate relative changes by comparing local labor markets that experience a more versus less severe recession. This variation, however, does not allow us to identify absolute changes in local economic activity following recessions (e.g., Nakamura and Steinsson 2014). Nonetheless, the persistent relative changes we find raise the concern that the capabilities of workers in some areas remain underutilized. This "direct effect" could lower aggregate output. At the same time, there could be an offsetting "indirect effect" if recessions reallocate employment to more productive areas. We examine this possibility through simple back-of-the-envelope calculations described in online Appendix B.5 and find no evidence of such productivity-enhancing reallocation. Fully assessing the impacts of persistent local labor market declines on aggregate output requires additional assumptions about the counterfactual evolution of economic activity in the absence of recessions, which we leave for future work.

Irrespective of the aggregate consequences of local labor market declines following recessions, our findings have important implications for labor market dynamism, the economic opportunities of workers and their children, and optimal policy

[56] This Monte Carlo exercise does not rule out other potential sources of bias when estimating equation (1), but we prefer to explore those issues using actual data.

[57] Online Appendix B.4 describes additional results which show that differences in the sample, time period, and level of geography do not explain why we find a persistent decrease in the employment-population ratio while the prior literature estimating BK VARs finds evidence of complete recovery.

responses. Our results show that recessions are followed by a sizable reallocation of employment across space. Local areas that experience more severe recessions see a persistent decline in employment across all sectors. At the same time, we find reductions in both in-migration and out-migration after local employment losses, which suggests that individuals are limited in their ability or willingness to move across areas to equilibrate shifts in labor demand. Moreover, the persistent decrease in local economic activity limits the opportunities available to both adults and children in these places. In response to these changes, investments in job creation and skill development could play an important role in boosting local economic activity. Such policies also could forestall the associated reduction in economic mobility for children (Stuart 2022). Currently, the vast majority of policy responses to recessions focus on short-term conditions. Our results imply that additional consideration should be paid to recessions' long-term consequences.

REFERENCES

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-Share Designs: Theory and Inference." *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.

Austin, Benjamin, Edward Glaeser, and Lawrence Summers. 2018. "Jobs for the Heartland: Place-Based Policies in 21st-Century America." *Brookings Papers on Economic Activity* Spring: 151–232.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi. 2020. "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure." *American Economic Review* 110 (10): 3139–83.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, Michigan: W. E. Upjohn Institute for Employment Research.

Bartik, Timothy J. 1993. "Who Benefits from Local Job Growth: Migrants or the Original Residents?" *Regional Studies* 27 (4): 297–311.

Bartik, Timothy J. 2015. "How Effects of Local Labor Demand Shocks Vary with the Initial Local Unemployment Rate." *Growth and Change* 46 (4): 529–57.

Bartik, Timothy J., Stephen C. Y. Biddle, Brad J. Hershbein, and Nathan D. Sotherland. 2019. "WholeData: Unsuppressed County Business Patterns Data: Version 1.0 [dataset]." Kalamazoo, Michigan: W. E. Upjohn Institute for Employment Research (accessed February 23, 2019).

Beaudry, Paul, David A. Green, and Ben M. Sand. 2018. "In Search of Labor Demand." *American Economic Review* 108 (9): 2714–57.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 1: 1–61.

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.

Cajner, Tomaz, John Coglianese, and Joshua Montes. 2021. "The Long-Lived Cyclicality of the Labor Force Participation Rate." Finance and Economics Discussion Series 2021-047.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3): 1405–54.

Charles, Kerwin Kofi, and Melvin Stephens Jr. 2013. "Employment, Wages, and Voter Turnout." *American Economic Journal: Applied Economics* 5 (4): 111–43.

Chetty, Raj, and Nathaniel Hendren. 2018a. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *Quarterly Journal of Economics* 133 (3): 1107–62.

Chetty, Raj, and Nathaniel Hendren. 2018b. "The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates." *Quarterly Journal of Economics* 133 (3): 1163–1228.

Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. "Regional Labor Market Adjustment in the United States: Trend and Cycle." *Review of Economics and Statistics* 99 (2): 243–57.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity* 2.

Dix-Carneiro, Rafael, and Brian K. Kovak. 2017. "Trade Liberalization and Regional Dynamics." *American Economic Review* 107 (10): 2908–46.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2020. "A Plucking Model of Business Cycles." Unpublished.

Federal Housing Finance Agency. 1975–2019. "Housing Price Index: Counties." Federal Housing Finance Agency. https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index-Datasets.aspx (accessed December 5, 2022).

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2021. "Place-Based Drivers of Mortality: Evidence from Migration." *American Economic Review* 111 (8): 2697–2735.

Freedman, Matthew. 2017. "Persistence in Industrial Policy Impacts: Evidence from Depression-Era Mississippi." *Journal of Urban Economics* 102: 34–51.

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79 (3): 733–72.

Garin, Andy, and Jonathan Rothbaum. 2022. "The Long-Run Impacts of Public Industrial Investment on Regional Development and Economic Mobility: Evidence from World War II." Unpublished.

Gathmann, Christina, Ines Helm, and Uta Schönberg. 2020. "Spillover Effects of Mass Layoffs." *Journal of the European Economic Association* 18 (1): 427–68.

Glaeser, Edward L., and Kristina Tobio. 2008. "The Rise of the Sunbelt." *Southern Economic Journal* 74 (3): 609–43.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8): 2586–2624.

Greenstone, Michael, and Adam Looney. 2010. *An Economic Strategy to Renew American Communities*. Washington, DC: Hamilton Project.

Hall, Robert E., 1992. "Discussion of 'Regional Evolutions' by Blanchard and Katz." *Brookings Papers on Economic Activity* 1: 62–65.

Hall, Robert E., and Marianna Kudlyak. 2020. "Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?" NBER Working Paper 27234.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (7): 1737–72.

Holmes, Thomas J., and John J. Stevens. 2002. "Geographic Concentration and Establishment Scale." *Review of Economics and Statistics* 84 (4): 682–90.

Hershbein, Brad, and Bryan A. Stuart. 2024. "Replication data for: The Evolution of Local Labor Markets After Recessions." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E184022V1.

Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller. 2012. "Who Suffers during Recessions?" *Journal of Economic Perspectives* 26 (3): 27–48.

Hurwicz, Leonid. 1950. "Least-Squares Bias in Time Series." In *Statistical Inference in Dynamic Models*, edited by Tjalling C. Koopmans, 365–83. New Haven, CT: Cowles Commission for Research in Economics.

Internal Revenue Service. 1993–2019. "Statistics on Income US Population Migration Data." United States Department of Treasury. https://www.irs.gov/statistics/soi-tax-stats-migration-data (accessed November 13, 2022).

Jacobson, Louis J., Robert J. LaLonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2012. "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey." *Demography* 49 (3): 1061–74.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration." *International Economic Review* 58 (1): 57–94.

Kilian, Lutz. 1998. "Small-Sample Confidence Intervals for Impulse Response Functions." *Review of Economics and Statistics* 80 (2): 218–30.

Kilian, Lutz. 1999. "Finite-Sample Properties of Percentile and Percentile-t Bootstrap Confidence Intervals for Impulse Responses." *Review of Economics and Statistics* 81 (4): 652–60.

Kilian, Lutz, and Helmut Lütkepohl. 2017. *Structural Vector Autoregressive Analysis. Themes in Modern Econometrics*. Cambridge, UK: Cambridge University Press.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury. 2020. "Sources of Displaced Workers' Long-Term Earnings Losses." *American Economic Review* 110 (10): 3231–66.

Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles. 2019. "IPUMS National Historical Geographic Information System: Version 14.0 [Database]" (accessed January 13, 2020).

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128 (4): 1687–1726.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2014. "Declining Migration within the US: The Role of the Labor Market." NBER Working Paper 20065.

Monras, Joan. 2020. "Economic Shocks and Internal Migration." Unpublished.

Moretti, Enrico. 2013. "Real Wage Inequality." *American Economic Journal: Applied Economics* 5 (1): 65–103.

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

National Cancer Institute. 1969–2019. "Surveillance, Epidemiology, and End Results US County Population Data." United States Department of Health and Human Services. https://seer.cancer.gov/popdata/download.html (accessed February 28, 2021).

Notowidigdo, Matthew J. 2020. "The Incidence of Local Labor Demand Shocks." *Journal of Labor Economics* 38 (3): 687–725.

Office of Management and Budget. 2003. *Bulletin No. 03-04*. Washington, DC: Executive Office of the President.

Pope, Arun Lloyd. 1990. "Biases of Estimators in Multivariate Non-Gaussian Autoregressions." *Journal of Time Series Analysis* 11 (3): 249–58.

Rinz, Kevin. 2022. "Did Timing Matter? Life Cycle Differences in Effects of Exposure to the Great Recession." *Journal of Labor Economics* 40 (3): 703–35.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. 2019. "IPUMS USA: Version 9.0 [dataset]" (access February 28, 2021).

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom. 2019. "Skewed Business Cycles." NBER Working Paper 26565.

Schmieder, Johannes F., Till von Wachter, and Joerg Heining. 2023. "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany." *American Economic Review* 113 (5): 1208–54.

Shaman, Paul, and Robert A. Stine. 1988. "The Bias of Autoregressive Coefficient Estimators." *Journal of the American Statistical Association* 83 (403): 842–48.

Stine, Robert A., and Paul Shaman. 1989. "A Fixed Point Characterization for Bias of Autoregressive Estimators." *Annals of Statistics* 17 (3): 1275–84.

Stock, James H., and Mark W. Watson. 2001. "Vector Autoregressions." *Journal of Economic Perspectives* 15 (4): 101–15.

Stuart, Bryan A. 2022. "The Long-Run Effects of Recessions on Education and Income." *American Economic Journal: Applied Economics* 14 (1): 42–74.

Temin, Peter. 1998. "Causes of American Business Cycles: An Essay in Economic Historiography." *Federal Reserve Bank of Boston Conference Proceedings* 42 (June): 37–64.

Tolbert, Charles M., and Molly Sizer. 1996. "US Commuting Zones and Labor Market Areas: A 1990 Update." Economic Research Service Staff Paper 9614.

US Bureau of Economic Analysis. 1969–2019a. "Personal Consumption Expenditures Price Index, Table 1.1.4, Line 2." United States Department of Commerce. https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=survey (accessed October 11, 2022).

US Bureau of Economic Analysis. 1969–2019b. "Regional Economic Accounts: CAGDP9 Real GDP in Chained Dollars by County and MSA, CAINC1 Annual Personal Income by County, CAINC5N Personal Income by Major Component of Earnings by NAICS Industry, CAINC5S Personal Income by Major Component of Earnings by SIC Industry, CAINC25N Total Full-Time and Part-Time Employment by NAICS Industry, CAINC25S Total Full-Time and Part-Time Employment by SIC Industry, CAINC35 Personal Current Transfer Receipts, SAEMP25N Annual Personal Income and Employment by State." United States Department of Commerce. https://apps.bea.gov/regional/downloadzip.cfm (accessed November 30, 2021).

US Bureau of Labor Statistics. 1969–2019. "Current Employment Statistics: United States, Total Nonfarm, Seasonally Adjusted, CES0000000001; State, Total Nonfarm, Seasonally Adjusted, SMSXX000000000000001; State, Total Nonfarm, Not Seasonally Adjusted, SMUXX000000000000001; United States, Historical Seasonally Adjusted, Various Industries, EES10000001, EES20000001, EES30000001, EES40000001, EES51000001, EES60000001, EES70000001, EES80000001." United States Department of Labor. https://data.bls.gov/cgi-bin/srgate (accessed January 14, 2023).

US Bureau of Labor Statistics. 1975–2019. "Quarterly Census of Employment and Wages, Single Files." United States Department of Labor. https://www.bls.gov/cew/downloadable-data-files.htm (accessed November 30, 2021).

US Bureau of Labor Statistics. 1976–2019. "Local Area Unemployment Statistics: Unemployment Rate, LAUSTXX0000000000003; Unemployment, LAUXX010000000000004; Employment, LAUXX010000000000005; Labor Force, LAUXX010000000000006." United States Department of Labor. https://download.bls.gov/pub/time.series/la/ (accessed August 29, 2022).

US Census Bureau. 1970–1994. "County Business Patterns. ICPSR 24722, 8464, 8441, 8442, 8142, 8348, 8360, 8433, 8665, 8883, 9198, 9381, 9711, 9740, 6030, 6382, 6488, 2343, 2280." ICPSR. https://www.icpsr.umich.edu/web/ICPSR/series/22 (accessed February 1, 2022).

US Census Bureau. 1995–2017. "County Business Patterns: Complete County Files." United States Department of Commerce. https://www.census.gov/programs-surveys/cbp/data/datasets.html (accessed February 14, 2021).

Wilson, Riley. 2022. "Moving to Economic Opportunity: The Migration Response to the Fracking Boom." *Journal of Human Resources* 57 (3): 918–55.

Wozniak, Abigail. 2010. "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources* 45 (4): 944–70.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127 (5): 2505–58.