

NBER WORKING PAPER SERIES

SHOULD I STAY OR SHOULD I GO?
THE RESPONSE OF LABOR MIGRATION TO ECONOMIC SHOCKS

Andrea Foschi
Christopher House
Christian Proebsting
Linda Tesar

Working Paper 33755
<http://www.nber.org/papers/w33755>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2025

We gratefully acknowledge financial support from the Michigan Institute for Teaching and Research in Economics (MITRE) and from the Brookings Institution. We also thank Brian Kovak, Abigail Wozniak and Jon Steinsson for their comments on an earlier draft of this paper. The views expressed here are those of the authors and do not necessarily reflect those of the Bank of Italy, the Eurosystem, or the National Bureau of Economic Research. Note: This is an earlier version of the paper prepared for the Spring 2025 Brookings Papers on Economic Activity (BPEA) Conference and the final version of this paper will be published in the Spring 2025 BPEA issue.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Andrea Foschi, Christopher House, Christian Proebsting, and Linda Tesar. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks
Andrea Foschi, Christopher House, Christian Proebsting, and Linda Tesar
NBER Working Paper No. 33755
May 2025
JEL No. E24, E32, F66, J61, R23

ABSTRACT

We examine the responsiveness of labor participation, unemployment and labor migration to exogenous variations in labor demand. Our empirical approach considers four instruments for regional labor demand commonly used in the literature. Empirically, we find that labor migration is a significant margin of adjustment for all our instruments. Following an increase in regional labor demand, the initial increase in employment is accounted for mainly through a reduction in unemployment. Over time however, net labor in-migration becomes the dominant factor contributing to increased regional employment. After 5 years, roughly 60 percent of the increase in employment is explained by the change in population. Responses of labor migration are strongest for individuals aged 20-35. Based on historical data back to the 1950s, we find no evidence of a decline in the elasticity of migration to changes in employment.

Andrea Foschi
Banca d'Italia
afoschi@umich.edu

Christopher House
University of Michigan
Department of Economics
and NBER
chouse@umich.edu

Christian Proebsting
KU Leuven
Department of Economics
Christian.Probsting@kuleuven.be

Linda Tesar
University of Michigan
Department of Economics
and NBER
ltesar@umich.edu

1 INTRODUCTION

In 2010, the population of Williston, North Dakota was 14,716 people – up from 12,512 a decade earlier. By 2020 the population in Williston had increased to 29,160 people. The reason for the dramatic surge in population was the discovery and development of oil fields in the Bakken Formation. Workers were drawn to Williston and other nearby towns because of high and rising wages and record low unemployment rates. The migration of labor has a long history in the United States. The California gold rush of the 1850s and the regional construction booms of the early 2000s are both well-known examples of workers moving to opportunity. In contrast, Detroit, Michigan – once one of the largest and wealthiest cities in the U.S. – has been shedding workers and residents for decades due to the declining fortunes of the American auto industry. According to the Census, in 1950, Detroit had a population of more than 1.8 million people. By the time of the 2020 decennial census however, Detroit’s population was only 640 thousand. The ebb and flow of workers across regions is clearly part of the dynamism of the U.S. labor market. In a well-known study of labor market dynamics in 1992, Blanchard and Katz (1992, hereafter: BK) concluded that the *dominant* equilibrating mechanism to a regional shift in labor demand was worker migration. That is, following an increase in labor demand, while local unemployment falls and labor participation rises, the main source of additional labor comes from workers moving to the region from elsewhere.

In contrast to the early study by BK, there is a growing sense among researchers that workers today are less willing to relocate from struggling locations to booming locations than were their predecessors.¹ Indeed, as Figure 1 shows, the rate of migration between U.S. states has declined over the last 40 years. The figure shows the gross cross-state migration rate (the average of in-migration plus out-migration divided by state population) since 1976. The figure is based on data from the Internal Revenue Service (IRS) that capture the year-on-year changes in address associated with tax returns. The figure supports the view that migration rates declined steadily from 1975 to the early 2000s, with some evidence that they have since stabilized at roughly 3 percent. The secular decline in labor mobility is considered to be an important contributor to the overall decline in U.S. labor market flexibility (Molloy et al., 2016).

The fact that migration rates have declined, however, does not necessarily mean that labor mobility has become less responsive to regional economic disturbances. In this paper we examine the joint responsiveness of labor migration, labor force participation and unemployment to plausibly

¹Jia et al. (2023) provide a comprehensive overview of recent research on changes in U.S. migration patterns.




Figure 1: GROSS CROSS-STATE MIGRATION RATE, 1976-2021

Note: The figure plots the gross cross-state migration rate, based on IRS data, computed as $\frac{1}{48} \sum_{i=1}^{48} 0.5 \frac{In-migrants_{i,t} + Out-migrants_{i,t}}{Non-migrants_{i,t} + Out-migrants_{i,t}}$, i.e. for each year we divide the average number of migrating tax returns by the number of all tax returns observed in t that originate from state i , and then average across states. Note that we exclude Alaska and Hawaii. A number of authors have pointed to inconsistencies and changes in methodology in the collection and production of the data starting from 2011-2012, when the IRS took over responsibility for providing these figures from the Census (see, e.g. DeWaard et al., 2022). These changes make the post-2012 data not comparable with the pre-2012 data, and likely explain the wild swings in 2014-2016, where the numbers are not reliable and are therefore greyed-out in the plot. Nonetheless, we believe post-2012 is still informative to have a sense of overall trends.

exogenous variations in labor demand. We abstract from long-run trends in migration flows between regions and instead focus on migration reactions to specific economic shocks. Our analysis allows us to gauge both the total regional response to labor demand shocks and the decomposition of the response. We base our analysis on four well-known instruments for labor demand at the regional level. Specifically, we consider an industry-share instrument (based on Bartik, 1993 and similar to Dao et al. (2017)), a military spending share instrument (based on Nakamura and Steinsson, 2014 and Auerbach et al., 2020), an import share instrument, i.e. the “China Shock” (based on Autor et al., 2021), and a regional housing instrument (based on Mian and Sufi, 2014). We consider three regional aggregation levels: states, commuting zones (CZs) and counties.

Our results suggest that labor mobility remains an important channel for labor adjustment despite the decline in gross migration rates. An increase in labor demand causes an equilibrium increase in employment. Initially, the increase in employment is accounted for primarily by a reduction in regional unemployment. According to our estimates, in the year of the labor demand shock,

roughly 80 percent of the increase in employment comes from lower local unemployment rates with the remaining 20 percent coming from elevated labor force participation and in-migration of workers. After five years however, the picture changes. Roughly 60 percent of the long-run change in employment is accounted for by the net migration of workers. The labor migration response is strongest for individuals in the 20-35 age range suggesting that much of the migration is due to younger workers seeking employment opportunities. The importance of the labor mobility channel is also more pronounced for urban, higher-income, and more college-educated counties, pointing to underlying differences in migration patterns by demographic characteristics.

Extending our dataset back to the 1950s allows us to examine changes in the migration elasticity over time. Using the industrial composition instrument, we show that the responses of employment and population at the 5-year horizon to labor demand shocks fluctuate over time. To the extent that researchers base their conclusions on different time periods, this could explain why they come to different conclusions about the importance of labor mobility. However, the *elasticity* of migration to labor-demand induced changes in employment (the ratio of the population response to the employment response) is approximately constant over the full sample.

We also further distinguish between jobs and employment, highlighting the importance of an additional adjustment margin, the jobs-to-employment ratio, that is especially relevant on impact at the state level and captures the importance of dual job-holding and commuting. Finally, we do not find significant differences in the role of mobility for positive versus negative shocks, but we decompose net migration into in-migration and out-migration, and into cross-state and within-state migration. We find in-migration is more important than out-migration, and cross-state mobility is more important than within-state mobility.

Our results speak to the large and growing literature on the extent of labor mobility in the United States. While BK argued that labor mobility was the dominant adjustment mechanism, subsequent studies found a smaller role for regional labor mobility (Dao et al., 2017; Fieldhouse et al., 2024). Studies focused on the Great Recession (e.g., Mian and Sufi, 2014 and Yagan, 2014) and on the effects of Chinese import penetration (e.g., Autor et al., 2021) find a smaller role for migration in equilibrating regional labor markets. The view that regional labor mobility has declined, has led some analysts to simply ignore migration as a possible adjustment mechanism (see for instance Nakamura and Steinsson, 2014; Beraja et al., 2019).

Not all recent research supports the view that labor migration has faded in importance however. For instance, Foote et al. (2019) study the effects of mass layoffs on regional employment and find that “the predominant adjustment mechanism following a mass layoff event is migra-

tion.” However, like other researchers, Foote et al. (2019) find that migration played a substantially smaller role during and after the Great Recession. In recent work, Foschi et al. (2023) show that while there was a temporary reduction in the responsiveness of migration in the years leading up to the Great Recession, migration elasticities during the crisis were comparable with the 1970-2000 period. Recent studies that focus on environmental disasters also find strong regional effects of migration as workers move out of distressed areas.

From a policy perspective, it is important to have a clear picture of how readily workers relocate in response to shocks. If the labor market is fluid and workers can relocate freely, then policy efforts can focus on addressing business cycles at the aggregate level. Alternatively, if workers are not able to move from hard-hit locations, place-based policies might be necessary to address regional downturns. Our paper provides estimates of both short-run and long-run migration elasticities, key factors in evaluating alternative policies.²

2 OVERVIEW OF RESEARCH DESIGN

We begin by discussing the basic econometric specifications used in our analysis. Our objective is to quantify the role that labor migration plays in the adjustment of employment following a shift in labor demand.³ We first present our decomposition of regional employment changes into changes in the unemployment rate, labor force participation rate, and regional population (i.e. migration). After discussing the data for these variables, we present OLS estimates on the historical co-movement between population and employment. Aware of the shortcomings of this setup, we then discuss our econometric approach to decompose labor-demand-driven changes in employment into changes in unemployment rates, participation rates and population.

²In recent work, O’Connor (2024) studies the intertemporal tradeoff between policies that support workers in hard-hit areas and providing incentives to relocate in the long-run.

³Note that we restrict our attention to changes in the extensive labor margin, that is changes in the number of workers employed. We do not study changes in the intensive labor margin—changes in the number of hours worked per person.

2.1 DECOMPOSITION OF EMPLOYMENT GROWTH

For any region i and time t , we can express total regional employment, $E_{i,t}$, as⁴

$$E_{i,t} = \underbrace{\left(1 - \frac{U_{i,t}}{LF_{i,t}}\right)}_{1-ur_{i,t}} \times \underbrace{\frac{LF_{i,t}}{POP_{i,t}}}_{LFP_{i,t}} \times POP_{i,t}.$$

where $U_{i,t}$ denotes the total number of people unemployed, $LF_{i,t} = E_{i,t} + U_{i,t}$ the labor force, and $POP_{i,t}$ the population. Let $ur_{i,t}$ denote the unemployment rate and let $LFP_{i,t}$ denote the labor force participation rate respectively. Taking logs and first differencing gives us the following relationship which holds by definition:⁵

$$\Delta \ln E_{i,t} = \Delta \ln(1 - ur_{i,t}) + \Delta \ln LFP_{i,t} + \Delta \ln POP_{i,t}. \quad (1)$$

This identity allows us to decompose the log growth rate in employment (the left hand side) into the log change in the employment rate (which approximately equals the negative of the log change of the unemployment rate), the log change in the labor force participation rate and the log change in population. In what follows, we interpret the change in population as net migration because changes in population due to births and deaths are not directly affected by cyclical labor market shocks.⁶

⁴Employment refers to the number of people residing in region i that are employed (anywhere). This is different from the number of jobs in region i . Employees might work more than one job, and employees might live in a region different from the location of their job, causing the number employed to differ from the number of jobs in a location. It is possible that shocks to labor demand could affect the incidence of people working multiple jobs and the incidence of commuting between regions. We return to the distinction between employment and jobs later in Section 3.1.6.

⁵The labor force participation rate is typically defined as the number of employed or unemployed over the non-institutionalized working-age population rather than the total population. Using that definition, would require including a fourth term reflecting the ratio of the non-institutionalized working-age population to the overall population.

⁶The population at date t is the population at date $t - 1$ times the gross birth rate $1 + b_{i,t}$, the gross survival rate $1 - d_{i,t}$, and the net migration rate $1 + nm_{i,t}$. While we are not aware of any studies linking the number of births to local business cycles, both Sullivan and Von Wachter (2009) as well as Pierce and Schott (2020) report somewhat higher mortality among workers facing earning losses. For instance, Pierce and Schott (2020) report that counties more exposed to import competition from China see a slight increase in “deaths of despair”, often associated with drug overdoses. But quantitatively, these effects are small with an interquartile shift in counties’ exposure being associated with a relative increase in mortality by 1 per 100,000. We therefore treat the log change in population as being equivalent to the net migration rate. To a first-order approximation, the log change in population satisfies

$$\Delta \ln POP_{i,t} \approx \bar{b} - \bar{d} + nm_{i,t}.$$

In our regression analysis, we include time fixed effects which reflect any nationwide trends in birth and death rates.

2.2 OVERVIEW OF LABOR MARKET AND POPULATION DATA

Here, we present an overview of the data used for our analysis. Details about the data associated with specific instruments are discussed in separate sections below.

2.2.1 LEVELS OF GEOGRAPHIC DISAGGREGATION

Our analysis examines labor market and population responses at three levels of aggregation: states, commuting zones (CZs) and counties. We drop Alaska and Hawaii, on the grounds that they do not share a border with any other states and thus should be expected to have unusual migration patterns; we also drop the District of Columbia.

The most disaggregated regional data we use is county-level data — an annual panel of $N = 3,049$ counties. CZs are collections of counties grouped together in such a way that interrelated economic and labor market activities are contained within their borders. CZs never cross county lines, but they may cross state lines (e.g. New York and New Jersey). We use the 1990 definition of CZs, which gives us $N = 721$ CZs.⁷ The most aggregated regional data we use are at the state level ($N = 48$).

2.2.2 LABOR MARKET DATA

For each county, we obtain data on employment, unemployment, and labor force from the Local Area Unemployment Statistics (LAUS) of the Bureau of Labor Statistics (BLS).⁸ Employment from LAUS measures the number of people employed as opposed to the number of jobs and follows a place-of-residence concept as opposed to place-of-work (i.e. employed people are counted based on where they live and not on where they work). Data is available at the county level starting in 1976.⁹ Annual figures are computed as averages of the January-to-December monthly figures.

⁷We use the crosswalk by Autor et al. (2013) to map counties into CZs. Note that all counties are included in a CZ. On average, each CZ includes 4 counties. The smallest CZ has just one county and the largest CZ has 19.

⁸Employment and unemployment figures in LAUS are estimated using the so-called Handbook method, which uses as inputs a combination of data from the Current Population Survey (CPS), the Current Employment Survey (CES), the Quarterly Census of Employment and Wages (QCEW), the American Community Survey (ACS), unemployment insurance (UI) claims counts, and population estimates from the U.S. Census Bureau. The methodology ensures that figures add-up to their state-level estimates, obtained separately with the CPS as the main input, and the CES and UI claims as secondary inputs for employment and unemployment estimates, respectively.

⁹Official data start in 1990. We obtained unofficial data for 1976-1990 upon request from the BLS. The BLS warns that pre- and post-1990 data might not be fully comparable. In addition, there were a number of changes in the estimation methodology for county data, occurring in 1985, 2005, and 2015. The last two are part of decennial program redesigns, whose updates are applied to the whole decade, so that the corresponding potential breaks would appear in 2000 and 2010, respectively. We keep all these years for the baseline analysis, but carry out robustness checks by excluding them from the sample and find that they do not affect results. Note that state-level figures are not impacted

In extensions of our baseline analysis, we integrate this employment measure with two additional ones: these are the employment figures from the Bureau of Economic Analysis (BEA) and from the State and Area Employment, Hours and Earnings Statistics contained in the BLS's CES. The latter are available at the state level dating back to 1945, which allows for a more historical extension of the analysis. The BEA data are available since 1969 and allow for finer disaggregation down to the county level. Unlike the LAUS, both these employment measures count the number of jobs and follow a place-of-work concept; for this reason, we will refer to them as "jobs" rather than "employment" throughout the paper.

2.2.3 POPULATION AND MIGRATION DATA

For each county, we take population data from the BEA.¹⁰ This is available at the county level since 1969, and at the state level since 1929. Population figures are reported as of July 1: to make them comparable to the annual averages of the labor market data, we use the following adjustment to construct the equivalent of a January-December average: $\overline{Pop}_{i,t} = 0.25Pop_{i,t-1} + 0.5Pop_{i,t} + 0.25Pop_{i,t+1}$. This adjustment assumes that population changes are equally spread across time between the observed dates. We use the change in total population in a region as our main measure of net migration. For some regressions we break down the BEA population data using two additional data sources. First, the Census also produces population data with a more detailed demographic breakdown (age, sex, and race), which is published by the Survey of Epidemiology and End Results (SEER) of the Cancer Institute and has the same time and geographic coverage as the BEA population data.¹¹ Second, the Internal Revenue Service (IRS) has data on tax return addresses, which is one of the inputs for the production of postcensal estimates, and results in a separate dataset with individual migration inflows and outflows by county for the 1990-2022 period and by state for the 1976-2022 period.¹² The IRS data allow us to decompose net migration

by these changes.

¹⁰This corresponds to the population data produced by the Census Bureau, consisting of decennial counts for Census years, and intercensal estimates for non-Census years. Intercensal estimates are produced using a form of interpolation to update postcensal estimates by correcting for the error of closure once a new Census year is reached. Postcensal estimates in turn are produced by separately estimating the components of population change (births, deaths, and domestic and international migration) and adding them up. Inputs for this estimation are vital statistics, data from the Social Security Administration's Numerical Identification File, Medicare enrollment, tax returns from the Internal Revenue Service (IRS), and the American Community Survey (ACS). The IRS tax returns data also form the basis for the IRS migration flow estimates described later.

¹¹See <https://www.nber.org/research/data/survey-epidemiology-and-end-results-seer-us-state-and-county-population-data-age-race-sex-hispanic>.

¹²Estimates of migration flows based on the IRS tax returns were produced by the Census and published by the IRS between 1990 and 2011-2012. For the 1976-1990 period at the state level, we rely on data provided by Molloy, Smith, and Wozniak (2011). Since 2011-2012, the migration figures are both produced and published by the IRS, which

into immigration and emigration and to calculate gross migration rates.¹³

2.3 REGIONAL EMPLOYMENT GROWTH AND REGIONAL NET MIGRATION

Consider a simple least squares regression of $\Delta \ln POP_{i,t}$ on $\Delta \ln E_{i,t}$ – two of the key variables we focus on in this paper. This regression will typically have a positive coefficient, so regions with high employment growth tend to be regions that experience net in-migration (i.e., $\Delta \ln POP_{i,t} > 0$). This could be because regions with high labor demand tend to have high wages and low unemployment and thus tend to attract workers from other parts of the country. On the other hand, a positive coefficient could also arise because workers move to a new area for personal reasons (e.g. for better weather, to be closer to family, etc.) and then seek work when they get there. Disentangling these two basic kinds of regional labor flows (labor-demand driven versus labor-supply driven) is the main focus of this paper. Nevertheless, it is instructive to consider the results of the simple OLS specification above as a starting point.

Here we consider the regression

$$\widehat{\Delta \ln POP}_{i,t+h} = \alpha_h + \beta_{h,t} \widehat{\Delta \ln E}_{i,t} + \varepsilon_{i,t+h} \quad (2)$$

for different horizons $h = 0, 1, 2, \dots$ and years t . The i regions are U.S. states (excluding Alaska and Hawaii). We denote the log difference in population between $t-1$ and $t+h$ by $\Delta \ln POP_{i,t+h}$. Note that we run this regression period by period, so that we estimate a different $\beta_{h,t}$ coefficient for each year and each horizon. Hats are used to indicate that the variables have been double-demeaned as follows (for any log-differenced variable $Y_{i,t}$):

$$\widehat{Y}_{i,t} = Y_{i,t} - Y_i - (Y_t - \bar{Y}) .$$

Here, $Y_i = \frac{1}{T} \sum_t Y_{i,t}$ is the long-run average for the state; $Y_t = \frac{1}{N} \sum_i \frac{Pop_i}{Pop} Y_{i,t}$ is a weighted year-by-year cross-sectional average; and $\bar{Y} = \frac{1}{T} \sum_t Y_t$ is the time-series average of the cross-sectional means. Double-demeaning removes long-run average differences as well as common cyclical vari-

also rolled out some methodological changes; these changes are likely the reason behind the unusual fluctuations in gross migration observed in 2014-2016, but, since they seemed to apply to both inflows and outflows similarly, the net migration figures are not as affected. Since the IRS data is also an input into the production of the intercensal population estimates, we conduct robustness checks by dropping the years after the new methodology is introduced – we find that the population response is the same.

¹³The American Community Survey provides another source of information about labor migration between states. As we noted in FHPT, problems with the collection of ACS data make it a less reliable data source for the study of migration than the IRS/Census data.

ations. We do this because many regions have persistently high (or low) unemployment rates and persistent in-(or out-) migration.

Figure 2a shows the estimated coefficient $\beta_{h,t}$ for $h = 0$ for each year over the 1945-2021 period. Data on employment by state is only available since 1976. To expand the analysis backwards, we also consider regressing population changes on changes in the number of jobs, for which longer time series are available. Figure 2b shows the same year-by-year coefficients for the five-year $h = 5$ cumulative population change. The average contemporaneous coefficient in Figure 2a is 0.26 over the full sample for jobs and 0.20 for employment. While there are some fluctuations around the mean, the relationship has been fairly steady over the past eighty years. In a given year, states that have employment or job growth that is one percent above average tend to have current net in-migration equal to roughly one fifth to one quarter of a percent of the population. The migration effect clearly becomes stronger over longer horizons (Figure 2b). At the $h = 5$ year horizon, the estimated coefficients are approximately 1.




Figure 2: COEFFICIENT OF NET MIGRATION ON JOBS AND EMPLOYMENT GROWTH, 1946-2021

Note: The figure plots the estimated coefficient β_h obtained from running (2) for $h = 0$ (left) and $h = 5$ (right) as a repeated cross-section year by year. The x-axis displays year t , so that with $h = 5$ panel (b) stops in 2016 but uses data up to 2021.

While these estimates are suggestive, without an instrument for labor demand, we cannot draw a causal inference from higher employment to migration inflows. In the sections that follow, we examine a number of different instruments for labor demand shocks to isolate the impact of changing employment opportunities on the incentive to migrate and confirm that the net migration response is consistent over time and across instruments.

2.4 ECONOMETRIC SPECIFICATION

We are interested in understanding the role that labor migration plays in how regions in the United States adjust to changes in employment driven by labor demand. Here we discuss our estimation strategy for a generic response to a shift in labor demand. We describe the specific labor demand instruments in greater detail in Sections 3.2.

Given a set of regional labor demand instruments $Z_{i,t}$, we run a series of horizon-specific regressions (i.e., local projections) for each variable of interest. Specifically, for each horizon $h = 0, 1, \dots, H$, and for any left-hand-side variable $Y_{i,t}$, we estimate the following regression:

$$Y_{i,t+h} = \alpha_{i,h}^Y + \alpha_{t,h}^Y + \beta_h^Y Z_{i,t} + \Gamma_h^Y X_{i,t} + \varepsilon_{i,t+h}^Y \quad (3)$$

The dependent variable $Y_{i,t+h}$ will be $\Delta \ln E_{i,t+h}$, $\Delta \ln(1 - ur_{i,t+h})$, $\Delta \ln LFP_{i,t+h}$ or $\Delta \ln POP_{i,t+h}$. Our baseline specification includes both horizon-specific time and region fixed effects. We discuss the control vectors $X_{i,t}$ below. Standard errors are clustered at the region level i , and regressions are weighted by a region's population share relative to the national population. The β coefficients give the estimated log change in employment, unemployment, labor force participation and population at time $t + h$ associated with a change in the instrument at date t .

In addition to reporting the β coefficients, we also report responses normalized by the change in employment. That is, for any left-hand-side variable $Y_{i,t}$ other than employment itself, we calculate¹⁴

$$\gamma_h^Y = \beta_h^Y / \beta_h^E. \quad (4)$$

These γ_h^Y coefficients reflect the fraction of the change in employment attributed to changes in Y . We are especially interested in γ_h^{POP} , the elasticity of population to employment. For instance, if a three percent increase in regional employment was achieved through a two percent increase in net migration and a one percent decrease in unemployment then we would have $\beta^E = 0.03$, $\beta^{POP} = 0.02$ and $\beta^{ur} = -0.01$ while $\gamma^{POP} = 0.66$ and $\gamma^{ur} = -0.33$, suggesting that migration accounts for two thirds of the employment response. It is important to note that this calcula-

¹⁴We use the Delta method to calculate the corresponding standard errors. To obtain the necessary estimates of the variance-covariance matrix, we stack the dataset, saturate it with a sample indicator, and run all regressions together (Angrist et al., 2023; Angrist and Hull, 2023). An alternative would be to estimate γ_h^Y and its standard errors directly by regressing $Y_{i,t+h}$ on $\Delta \ln E_{i,t+h}$ using $Z_{i,t}$ as an instrument for $\Delta \ln E_{i,t+h}$. This leads to the same point estimates, but somewhat different standard errors. In our setup, the standard errors from this 2SLS approach are always smaller.

tion masks any potential composition effects of migration. For instance, if in-migrants have relatively higher labor participation rates then some of the measured change in LFP or ur will be attributable to migration. This could be particularly true if in-migrants move only once they have secured a job in their new location.¹⁵

Given our decomposition (1), the estimated contributions should exactly sum to 1 for each horizon h regardless of the econometric specification and provided that the right-hand-side variables in (3) are the same for each independent variable Y :

$$-\hat{\gamma}_h^{ur} + \hat{\gamma}_h^{LFP} + \hat{\gamma}_h^{POP} = 1.$$

Of course, the β coefficients and the γ coefficients are simply the first-stage and second-stage estimates from a 2SLS specification. However, for our purposes, they convey important information about the changing nature of regional labor migration over time. The conventional wisdom is that migration has become less responsive to economic shocks in the US since the early 1990s. Below we will demonstrate that the elasticity of labor migration to shifts in employment demand has remained more or less constant, but that the persistence of the employment response fell in the 1990s. This will translate into declining β coefficients and roughly stable γ coefficients over time.

The main challenge in conducting this analysis is the identification of exogenous changes in the demand for labor that differentially affect regions across the United States; that is, in order to say that a shift in labor demand in a given locality causes local employment to rise, we need an instrument for local labor demand. We consider four widely-used proxy variables for a demand shock: regional differences in industry demand (an industry Bartik), regional differences in government spending, regional differences in Chinese import penetration and regional differences in house price changes. These proxies for labor demand are instrumented using shift-share instruments and are discussed in detail below.

3 ESTIMATES OF THE MIGRATION ELASTICITY

In this section we present our estimates of the migration elasticity, γ^{POP} . Section 3.1 presents results for the traditional Bartik instrument which takes shifts in industrial composition as the

¹⁵Generally, if migrants have a higher employment rate than the average population, γ^{POP} will underestimate the true contribution of migration. For example, suppose a state's labor force participation rate is 60% and its unemployment rate is 5%. If an increase in labor demand by one worker is met by a migrant, and that individual arrives alone, the regression would assign only 60% of the adjustment to the population margin.

regional labor demand instrument. We discuss the results in this section at some length before moving on to the other instruments. Section 3.2 considers other common instruments for labor demand.

3.1 INDUSTRY-MIX INSTRUMENT

We begin by considering the industry-mix instrument originally proposed by Bartik (Bartik, 1993) and later used by Dao et al. (2017) and Amior and Manning (2018) in their analyses of migration as a mechanism of labor market adjustment. This instrument is obtained by constructing, for each region, the growth of jobs that would have occurred given the historical industrial composition of jobs in the region, had each sector grown at the national growth rate. For region i at time t we calculate

$$Z_{i,t}^{IND} = \sum_{j=1}^J \bar{s}_{i,t-4:t-1}^j \frac{\Delta \mathcal{E}_{t,-i}^j}{\mathcal{E}_{t-1,-i}^j}.$$

In this equation, $\mathcal{E}_{t,-i}^j$ is the national number of jobs in industry j at time t , excluding region i ; and $\bar{s}_{i,t-4:t-1}^j$ is the share of industry j in total jobs, in region i , averaged over the previous four years ($t - 4$ to $t - 1$).¹⁶ The intuition for the Bartik industry instrument is that aggregate, national fluctuations in an industry will be associated with changes in labor demand for regions that specialize in that industry. For example, an increase in the national number of automobile jobs causes a disproportionate increase in labor demand in Michigan relative to other states that are less specialized in auto production. Note that the control vector (i.e., $X_{i,t}$ in regression (3)) includes a lag of the instrument.¹⁷

The industry-mix instrument has been widely used as a proxy for local labor demand shocks (see among many other Bound and Holzer, 2000; Beaudry et al., 2014; Notowidigdo, 2020). A recent literature in applied microeconometrics has helped clarify under what conditions these shift-share instruments yield consistent estimates (Adao et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022b). Most of the literature focuses on the single-cross section case, which raises concerns that instruments might be correlated with alternative explanations. The panel

¹⁶We also consider alternative ways of constructing the shares: we first average them over the previous 9 years (so that the shares for, e.g., 1990 are constructed as the average of the 1981-1989 shares); we then average them over the first 5 years of the previous 9 years, skipping the last 4 years (so that the shares for 1990 are constructed as the average of the 1981-1985 shares). The results are robust.

¹⁷We also ran regressions including a lag of the left-hand-side variable. The results are similar, but the resulting γ coefficients of these regressions do not exactly add up to 1 because the right-hand-side variables are not the same for each independent variable.

dimension is helpful in this regard because the large number of shocks in the time dimension reduces concerns of spurious correlation. The panel dimension also allows us to include time fixed effects and state fixed effects, which soak up a number of confounding factors and deal with time-varying means of the shifts (Borusyak et al., 2025).

Borusyak et al. (2025) point out that exogeneity of the time series variation in the shifts is sufficient for consistent estimation. Since employment growth at the national level is an equilibrium outcome that is affected by both labor demand and labor supply, it is hard to argue that the time-series variation in the shifts is exogenous. The main concern is that the instrument picks up local labor supply shocks. For instance, if workers move to Arizona in response to more widespread availability of air conditioning, then the industries concentrated in Arizona will grow more quickly, and the industry-mix instrument might capture labor supply shocks. Leaving out the state’s employment when calculating the shifts in the instrument partially addresses this concern, unless the labor supply shocks are correlated with states’ industry compositions (e.g. if workers also moved to Texas, and Arizona and Texas have similar industries).

Further below, we discuss that labor supply and labor demand shocks should move unemployment rates in opposite direction. Comparing our OLS estimates to those from using the industry mix instrument suggests that the instrument does a better job at isolating labor demand shocks, though some correlation with local labor supply shifts remains possible. In Section 3.2, we therefore look at alternative labor demand instruments and generally find similar results.

DATA We construct the Bartik by combining data from the Quarterly Census of Employment and Wages (QCEW), published by the BLS, to construct the shifts, and from the Census’ County Business Patterns, in the version provided by Eckert et al. (2021), to construct the shares.¹⁸ As in the BEA and the CES, the employment concept in the QCEW and the CBP is a count of the number of jobs by place of work. The advantage of using QCEW and CBP data for the construction of the Bartik is the much richer level of industry detail that they allow: in creating the Bartik, we use data at the highest available level of industry detail, which means 6-digits when using NAICS and 4-digits when using SIC; if that degree of detail is not available for an industry, we take data from the next higher level of aggregation. This means that we can rely on approximately 1000 industries

¹⁸While the QCEW also has data at the county level, we only use it for the computation of national changes in industry employment, and rely on the Eckert et al. (2021) data to compute employment shares at the county level. A key advantage of the Eckert et al. (2021) dataset is overcoming the problem of suppression: this is common to both the QCEW and the original CBP data when looking at very detailed industries at the county level, for confidentiality reasons. As a result, we believe that shares computed from the Eckert et al. (2021) version of the CBP data give a more complete picture of the detailed industry structure of counties.

overall.¹⁹ We build CZ- and state-level versions of the Bartik as employment-weighted averages of the counties they contain.

3.1.1 BASELINE RESULTS

Figure 3 shows the estimated local projections for employment, the unemployment rate, the labor force participation rate, and the net migration rate in response to a one unit increase in the Bartik industry composition instrument. That is, the figure plots the β_h coefficients from equation (3) estimated on state-level data. The construction of the instrument suggests that the predicted change in employment should be roughly 1 percent.²⁰ The estimated response is slightly less than 1 percent on impact ($h = 0$) though the overall change in employment rises over time. The predicted change in state employment is roughly 2.5 percent five years after the initial innovation in the instrument. Unemployment falls on impact and remains almost 1 percent below its initial value at all horizons. There is also a modest estimated change in labor force participation.

Our main coefficient of interest is the migration coefficient β_h^{POP} plotted in the lower right panel. The estimates show that there is essentially no increase in population on impact. Instead, state population rises steadily following the innovation. After five years, state population is predicted to increase by more than 1.5 percent.

Figure 4 shows the estimates for the γ coefficients associated with the β coefficients in Figure 3. The shaded regions reflect one standard deviation error bands. Because the estimates must sum to one, we can interpret the point estimates as shares of the predicted employment response. We see that on impact, roughly 80 percent of the increased employment is accounted for by a reduction in state unemployment. In-migration and increased labor participation account for the remaining 20 percent. As time passes however, these shares switch. At the five-year horizon, more than 60 percent of the estimated increase in employment is due to increased state population

¹⁹For the shift, using the QCEW, we rely on SIC data until 1997, the last year when it is available; starting from 1998, we use NAICS data. Since the BLS reconstructed the NAICS going back to 1990, we compute the 1997-1998 change in employment using the reconstructed 1997 NAICS data. This ensures there is no single period in which, between one year and the previous one, we are computing differences in employment between different industry classifications, thus preventing unexpected jumps. Similarly, for the shares, using CBP data, we rely on the concordance tables constructed by Eckert et al. (2021) to convert old industry codes to new industry codes whenever we get to a year when a new classification system is introduced. For instance, in 1997, when the shift variable is based on SIC, we compute the average shares over 1993-1997 using the native SIC shares used in those years; however, in 1998, when the shift variable switches to NAICS, we convert SIC codes to NAICS codes for 1994-1997 using the concordance tables, append those to 1998, and finally compute the average shares. This ensures that we have always harmonized the classification before computing averages.

²⁰More accurately, because the Bartik is based on a measure of jobs, the predicted change in jobs equals one, and the predicted change in employment will be less than 1.




Figure 3: RESPONSE TO BARTIK INSTRUMENT AT THE STATE LEVEL

Note: The figure plots the estimated β_h coefficients obtained from running (3) for each labor market variable at different horizons h (x-axis) using the state-level dataset. The sample period is from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

while only 20-25 percent comes from reduced unemployment.

It is worth comparing the results to the literature: Based on a VAR identified by the assumption that current shocks to employment growth are driven by labor demand only, BK find a substantially larger response of migration, accounting for half the employment response upon impact and all of it after five years. As pointed out by Dao et al. (2017), the identifying assumption in BK is likely to be violated as both labor demand and labor supply shocks shift employment. For instance, labor demand shocks that draw workers from other states also influence labor supply in the states those workers leave, causing population, employment and unemployment rate to fall. Alternatively, a state that observes a particularly large share of its workforce retiring might face labor shortages that drive down employment (and the unemployment rate), but pull in workers from other states. The correlation between population and employment induced by these labor supply shocks is therefore ambiguous, depending on whether the labor supply shock is internal to the state or external. But in both cases, labor supply shocks induce a positive correlation between unemployment rates and employment growth, in contrast to a negative correlation in the




Figure 4: RATIO BETWEEN EACH LABOR MARKET RESPONSE AND THE EMPLOYMENT RESPONSE

Note: The figure plots the overlapped γ_h coefficients defined in (4) for each labor market variable at different horizons h (x-axis) using the state-level dataset. The sample period is from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

presence of labor demand shocks.

To the extent that the industry-mix instrument better isolates shifts in labor demand compared to an OLS regression, we therefore expect a larger share of any increase in employment to be accounted for by a fall in the unemployment rate. This is indeed the case: The unemployment margin accounts for 81 percent of the impact employment response when using the industry-mix instrument, but only 32 percent when using OLS. Dao et al. (2017) also use a shift-share design to instrument for labor demand shocks and find comparable results for the impact responses: the unemployment margin dominates the population margin in the short run (79 percent vs. 11 percent, see their Table 1), in line with our γ estimates in Figure 4.

3.1.2 HAS MIGRATION BECOME LESS RESPONSIVE OVER TIME?

An important question is whether the estimates of labor migration elasticities have changed over time. As we noted previously, several recent papers seem to find reduced levels of labor migration. To get at this issue, we consider a sequence of rolling regressions in which we estimate equation (3) on shorter sub-samples of the original data. The rolling regressions allow us to see how the estimates change with different segments of data and show us whether the responsiveness of net migration to local economic conditions has changed over time. Specifically, we run the following

regression

$$\hat{Y}_{i,t+h} = \alpha_h^Y + \beta_h^Y \hat{Z}_{i,t} + \Gamma_h^Y \hat{X}_{i,t} + \varepsilon_{i,t+h}^Y, \quad (5)$$

where $Y_{i,t}$ is the cumulative log-change in jobs (based on the BLS-CES) and population (based on the Census). Here, $Z_{i,t}$ is a composite Bartik: from 1976 onwards, it corresponds to our baseline 4- and 6-digit Bartik constructed from the QCEW and the CBP; for pre-1976, it corresponds to a 1-digit Bartik, constructed analogously using data from the CES. The control X includes a lag of the instrument. Recall that hats are used to denote double-demeaned variables. Here, we split the sample in two for this double-demeaning: pre-1976 values are demeaned using pre-1976 means for Y_i and $\bar{Y}$; post-1976 values are instead demeaned using post-1976 means. This allows us to control for long-run trends in a way that is analogous to our fixed-effects specification (which is less suited for short 10-year windows) but does not impose the strong assumption of a single trend for each state since the 1940s up until today. We focus on two dependent variables: jobs and population.²¹ Finally, the regression is population-weighted and standard errors are clustered by state.

The upper panels of Figure 5 shows the estimated β_5 coefficients on employment and population. Each dot in the figure is an estimated coefficient for a ten-year rolling window (the dots are plotted over the midpoint in each window, so, for instance, the coefficient for the 1980-1989 interval is plotted on 1984.5). Based on our estimates, both state population and state employment increase five years after the innovation in instrument. The effect of the labor demand shock changes considerably over time and, according to the figure, has increased over the sample. Importantly, the ratio of the population response to the employment response, the γ^{POP} coefficient, has remained roughly constant over the past seven decades, as shown by the nearly horizontal best-fit line in Figure 5c.²² Most of the variation in the γ^{POP} coefficient arises during periods when

²¹A few issues prevent us from exactly replicating the baseline decomposition from Figures 3-4 with pre-1976 historical data: the 4/6-digit Bartik is only available starting 1976; the data on employment (based on number of employees by residence) is only available starting in 1976 (hence why we use jobs here); the data on unemployment is similarly only available starting in 1976; consequently, labor force participation (which requires number of employed plus number of unemployed) is also only available starting in 1976. As a result, population is the only variable from Figures 3-4 that has consistent data both pre- and post-1976. Employment could be replaced by jobs (as we do here), unemployment may be replaced by a reconstructed series (e.g. Fieldhouse et al. (2024)), but labor force participation would still not be available.

²²The attentive reader might notice that the γ value is below the value of 0.61 reported in Table 1. This is partly due to the longer sample. The γ has slightly increased over time. Partly this is due to us considering the response of the number of jobs in the denominator rather than the employment response. See Section 3.1.6 for a discussion on this distinction between jobs and employment.




Figure 5: HISTORICAL JOBS AND POPULATION RESPONSES TO BARTIK INSTRUMENT AT THE STATE LEVEL, 10-YEAR ROLLING WINDOWS

Note: The figure plots, at the top, the estimates of the β_h coefficients obtained by running (5) at horizon $h = 5$ using double-demeaned state-level data over 10-year rolling windows (full sample period: 1946 - 2016); at the bottom, the γ_h coefficient defined in (4) is obtained by taking the ratio of the resulting β_h coefficients for population and jobs. Note that the coefficients are plotted against the midpoint of the rolling window, so the coefficient for 2010.5 refers to the 2007-2016 sample, and with $h = 5$, this means using data up to 2021. Bars represent 90% confidence intervals. The red line is a weighted least square fit, where the weights are the inverse of the standard errors.

employment responses were historically low and estimated with less precision, leading to noisier estimates of the γ^{POP} .

This analysis places recent research on labor mobility into a historical context. For example, comparing the original BK sample (1976–1990) with subsequent data, Dao et al. (2017) conclude that the “migration sensitivity to regional shocks has been strongly decreasing since the 1990s.” This is consistent with the β^{POP} coefficient depicted in Figure 5b: for the 1976–1990 samples (midpoints 1980.5–1985.5), the β^{POP} coefficient remains consistently above 2, whereas for 1990–2007 (midpoints of 1994.5 to 2002.5), it hovers mostly around 0.5. Therefore, focusing on this snapshot gives the impression that labor mobility indeed has plummeted. However, most of the drop in the population response stems from a less persistent employment response. The decline in the elasticity of migration to employment changes depicted in 5c is substantially smaller—and by the 2000s it has fully recovered. Viewed from the perspective of history, there is therefore little evidence that the contribution of migration to labor adjustment has diminished.

3.1.3 MIGRATION RESPONSES FOR SMALLER REGIONS: COMMUTING ZONES AND COUNTIES

We can also use the industry composition instrument to examine labor migration between smaller geographical areas. Here we consider estimated labor migration responses for commuting zones (CZs) and counties. Table 1 presents coefficient estimates for the $h = 0$ and $h = 5$ horizons. The top panel of the table shows the β coefficients for each of the components of the employment decomposition while the lower panel shows the γ coefficients—the elasticities with respect to employment. The top row shows the state-level estimates — these correspond to the local projections in Figure 3. The second and third rows show estimates for CZs and for counties.

Focusing first on the β coefficients, the estimates for the impact period are similar across the different regions. Notably, there is little to no net migration response in the impact period regardless of the level of aggregation. Interestingly, when we look at the longer-run reactions ($h = 5$) it seems that the estimates are systematically smaller as we consider smaller regions. At the state-level, the estimated change in employment after five years is 2.6 percent. However, the 5-year estimate for CZs is only 1.7 percent and the estimate for counties is less than 1.2 percent. This pattern continues to hold for the other variables. After five years, state population has increased by 1.6 percent while CZ population has increased by 1 percent and county population by 0.6 percent.²³ Nevertheless, for all regional aggregates, net migration becomes the dominant adjustment

²³These smaller estimates at the CZ and county level could be due to measurement error in the industry-mix variable. The exposure shares at the county level are partially based on imputed values because the original industry-

mechanism over time. The table also highlights a key distinction between the β and γ coefficients, echoing our findings on migration responses over time. While net migration appears less responsive to labor demand shifts at more granular levels—reflected in smaller β coefficients—this weaker response is almost entirely attributable to the muted employment response. The elasticities with respect to employment (γ) are remarkably consistent across all levels of aggregation. At $h = 0$, the γ coefficient on migration ranges from 0.11 to 0.14; at $h = 5$ the coefficient increases to a range of 0.54 to 0.61.²⁴

Although the differences are not large, the point estimates suggest that migration elasticities are somewhat smaller at finer geographic levels (see also Foschi et al., 2023, for a similar finding). Intuitively, one might expect workers to be *more* mobile between counties than between states. After all, our county-level regressions capture both cross-state moves and within-state relocations.²⁵ Using IRS data, we separate county-level migration flows into cross-state and within-state net migration flows and re-run our local projections on each separately. While, unconditionally, 56 percent of county-level migration flows occur within state borders, within-state moves account for only a quarter of the migration response to labor demand shocks. This does not fully resolve the puzzling ranking of migration elasticities across levels of aggregation, but it mitigates the concern. While overall mobility is indeed higher at the county level, much of it appears unrelated to labor market conditions. Consequently, we should not necessarily expect migration elasticities to employment changes to be much higher at the county level than at the state level, but the fact that the point estimates are even lower remains puzzling.

3.1.4 MIGRATION RESPONSE BY DEMOGRAPHIC AND COUNTY CHARACTERISTICS

Up to this point we have relied on total population as the measure of migration flows. The Census data allow us to look at the change in population in response to employment shocks for different demographic groups. Figure 6 shows the coefficient on the change in population to state-level labor demand shocks at the five-year horizon by age. The responses show that individuals between the ages of 25 and 39 have the largest response to the labor demand shock and are therefore the

county level data is suppressed for confidentiality reasons (Eckert et al., 2020). Note that this measurement error cancels out when calculating the γ 's. We also found that labor demand changes are less persistent at a more granular level, which would also contribute to a smaller response.

²⁴These elasticities at the 5-year horizon are consistent with those estimated by Amior and Manning (2018) who use a similar framework that relies on the Bartik instrument. Based on decennial census data at the commuting zone level from 1950 to 2010, they estimate that a 10 percent decline in local employment leads to a 7 percent decrease in population over a decade.

²⁵Wilson (2022) even finds that country-to-country migration drops off discretely at state borders.

		Short-run ($h = 0$)				Long-run ($h = 5$)			
		Empl	ur	Lfp	Pop	Empl	ur	Lfp	Pop
β	State	0.69 (0.14)	-0.57 (0.09)	0.03 (0.07)	0.10 (0.04)	2.60 (0.32)	-0.62 (0.17)	0.39 (0.14)	1.59 (0.23)
	CZ	0.95 (0.06)	-0.53 (0.04)	0.31 (0.04)	0.11 (0.02)	1.73 (0.17)	-0.15 (0.07)	0.58 (0.09)	0.99 (0.09)
	County	0.63 (0.03)	-0.34 (0.02)	0.22 (0.03)	0.07 (0.01)	1.15 (0.08)	-0.07 (0.03)	0.47 (0.06)	0.62 (0.05)
γ	State		-0.82 (0.22)	0.04 (0.10)	0.14 (0.07)		-0.24 (0.07)	0.15 (0.06)	0.61 (0.11)
	CZ		-0.56 (0.06)	0.33 (0.05)	0.11 (0.02)		-0.09 (0.04)	0.34 (0.06)	0.57 (0.08)
	County		-0.54 (0.04)	0.35 (0.04)	0.11 (0.01)		-0.06 (0.03)	0.40 (0.06)	0.54 (0.06)

Table 1: LABOR MARKET RESPONSES AT DIFFERENT LEVELS OF AGGREGATION

Note: The top half of the table displays the estimates of the β_h coefficients obtained by running (3) for each of the dependent variables listed at the top, with $h = 0$ on the left part and $h = 5$ for the right part. The bottom half displays the γ coefficients, defined as the ratio between the β of the dependent variable of interest and the employment β^E . The sample period is from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Standard errors (in parentheses) for the β 's and γ 's are clustered at the level of the region. All regressions are weighted by population.

most likely to relocate.²⁶ We find no statistical distinction between men and women, or between races, although the standard deviations of the three race categories are large, making it difficult to draw sharp conclusions (these results are not shown).

Although we do not have micro-data to exactly capture the characteristics of movers, we collect data on county characteristics to assess potential differences in labor market responses associated with features of a region or its demography. In particular, we focus on differences between urban and rural counties; higher- and lower-income counties; and more- and less-educated counties.²⁷ For each pair of characteristics, we introduce an indicator variable into our estimation equation that differentiates between the two. We code a county as urban if it is either metropolitan, non-metropolitan but adjacent to a metropolitan county, or non-metropolitan with an urban population, according to the categorization provided in the 2013 Rural-Urban Continuum Codes (RUCC) by the Department of Agriculture (USDA). We code a county as higher-income if its income-per-

²⁶Kaplan and Schulhofer-Wohl (2017) and Molloy et al. (2017) examine the migration propensities of different age groups and the implications for labor flows. In general, their findings suggest that, to the extent there is declining labor mobility, it is not associated with the aging of the population.

²⁷We also look at income inequality, as measured by the Gini coefficient, but find that it is not associated with any differences in responses.




Figure 6: Population response by age

Note: The figure plots the estimated β_h coefficients obtained from running (3) for each labor market variable at horizons $h = 0$ (left) and $h = 5$ (right) using the state-level dataset. The sample period is from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Bars denote 90% confidence intervals.

capita (provided by the BEA) is in the top quintile for that year. Finally, we code a county as more educated if its share of adults completing at least some college (provided by the USDA) is in the top quintile for that year.

The results are presented in Table 2. Overall, labor mobility emerges as a more important adjustment mechanism for urban, richer, and more educated. For urban-vs-rural counties, this occurs because the employment response is similar, but the population response is stronger in urban than in rural counties (for reference, the corresponding γ coefficients at $h = 5$ for population relative to employment are 0.55 for urban and 0.42 for rural counties). For higher-income counties, this results from both a weaker employment response and a stronger population response than lower-income countries (the corresponding γ coefficients are 0.64 for higher-income and 0.45 for lower-income counties). Finally, for education, the employment response is again similar, but the population response is stronger for more educated than for less educated counties (the γ coefficients are 0.62 for more educated and 0.49 for less educated counties).

3.1.5 ASYMMETRIES IN MIGRATION RESPONSES

We modify our baseline regression to examine two types of asymmetries: (i) differences in the response to positive versus negative shocks and (ii) differences in in-migration versus out-migration. As in the previous section, we conduct the analysis at the state level.

Table 2: LABOR MARKET RESPONSES BY COUNTY CHARACTERISTICS

	Short-run ($h = 0$)				Long-run ($h = 5$)			
	Empl	ur	Lfp	Pop	Empl	ur	Lfp	Pop
Urban	0.69 (0.03)	-0.34 (0.02)	0.23 (0.03)	0.07 (0.01)	1.15 (0.08)	-0.06 (0.03)	0.46 (0.06)	0.63 (0.05)
Rural	0.45 (0.03)	-0.29 (0.01)	0.12 (0.03)	0.05 (0.01)	1.10 (0.08)	-0.15 (0.03)	0.48 (0.06)	0.46 (0.03)
Higher income	0.66 (0.04)	-0.31 (0.02)	0.26 (0.03)	0.09 (0.01)	1.06 (0.10)	-0.00 (0.03)	0.37 (0.06)	0.68 (0.06)
Lower income	0.61 (0.03)	-0.37 (0.02)	0.19 (0.02)	0.05 (0.01)	1.24 (0.08)	-0.13 (0.03)	0.56 (0.05)	0.56 (0.04)
Higher education	0.64 (0.04)	-0.29 (0.02)	0.26 (0.03)	0.09 (0.01)	1.18 (0.11)	-0.06 (0.03)	0.40 (0.08)	0.73 (0.07)
Lower education	0.63 (0.03)	-0.38 (0.02)	0.20 (0.02)	0.06 (0.01)	1.13 (0.07)	-0.07 (0.03)	0.51 (0.05)	0.55 (0.04)

Note: The top half of the table displays the estimates of the β_h coefficients obtained by running (3) for each of the dependent variables listed at the top, with $h = 0$ on the left part and $h = 5$ for the right part, and distinguishing between different county characteristics. Standard errors (in parentheses) are clustered at the level of the region. All regressions are weighted by population.

To distinguish between positive and negative shocks, we introduce an indicator variable into our estimation equation that differentiates between the two. We allow all model coefficients to vary based on the sign of the shock. Figure 7 presents the estimates. In the short run, we observe a stark asymmetry: the employment effect of a positive shock is less than half as large as that of a negative shock, with most of the additional response to negative shocks occurring through higher unemployment. This finding aligns with theories emphasizing search and matching frictions in labor markets, where the firing margin is more responsive than the hiring margin, as well as models that assume downward wage rigidity. In the long run, however, employment responses to positive and negative shocks converge, with net migration accounting for most of the adjustment in both cases.

Next, we decompose population changes into inflows and outflows using IRS migration data. Figure 8 shows that employment-driven migration is primarily driven by variation in in-migration, rather than out-migration. This pattern holds across all time horizons and is consistent with Monras (2018) who documents similar effects for the period of the Great Recession. Additional results (not reported here) suggest that this pattern is similar for both positive and negative shocks. This implies that when a state experiences a negative shock, the primary adjustment mechanism is

not an exodus of workers but rather a decline in the number of new arrivals. This reduction in labor supply helps stabilize the local labor market by lowering unemployment among those who remain.²⁸




Figure 7: Labor market responses to positive and negative labor demand shocks

Note: The figure plots the estimated β_h coefficients obtained from running (3) distinguishing between positive (left part) and negative (right part) shocks. Both the impact response ($h = 0$) and the five-year response ($h = 5$) are depicted for each labor market variable. The sample is at the state level and runs from 1976 till 2016, with projections going up to 2021 (for $h = 5$). The bars indicate 90% confidence intervals.

3.1.6 DISTINGUISHING BETWEEN JOBS AND EMPLOYMENT

Our baseline decomposition (6) breaks down employment growth into changes in the unemployment rate, labor force participation rate, and population growth. However, an additional margin of adjustment to labor demand shifts is the ratio of jobs to employment. When labor demand increases in region i , workers may take on multiple jobs, or some new jobs may be filled by commuters from outside region i . Incorporating jobs into the decomposition yields

$$\Delta \ln \mathcal{E}_{i,t} = \Delta \ln \frac{\mathcal{E}_{i,t}}{E_{i,t}} + \Delta \ln(1 - ur_{i,t}) + \Delta \ln LFP_{i,t} + \Delta \ln POP_{i,t}, \quad (6)$$

where $\mathcal{E}_{i,t}$ denotes jobs in region i at time t . Table 3 presents the β and γ estimates from re-estimating our local projections (3) at the state, CZ and county level but now including estimates for $\beta_h^{\mathcal{E}}$ and $\gamma_h^{\mathcal{E}}$.

²⁸ Monras (2018) rationalizes this finding in a structural model where households first decide whether to move or not and in a second step, choose the destination based on labor market conditions. This makes households less sensitive to economic conditions at home, but more sensitive to destination conditions conditional on moving.




Figure 8: In- and out-migration response to Bartik instrument at the state level

Note: The figure plots the estimated β_h coefficients obtained from running (3) for each migration measure at different horizons h (x-axis). The coefficient for out-migration has a flipped sign for readability and comparability with the net migration and out-migration responses. The sample is at the state level and runs from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

In our baseline, we observed that state-level employment initially rises by less than the 1 percent predicted increase in jobs (see Table 1). In contrast, the number of state-level jobs rises by 1 percent, consistent with the predicted increase. Thus, the number of jobs in region i rises more than the number of employed workers does. Quantitatively, the job-to-employment margin (column ‘je’) accounts for nearly one-third of short-run labor market adjustment at the state level. While the data does not distinguish between an increase in dual job-holdings and greater commuting, the near-zero job-to-employment margin at the CZ level suggests commuting plays a strong role, even across state borders (e.g., New Jersey and New York or Kansas and Missouri). Over time, this margin becomes less significant at the state level. In either case, in the longer run, migration remains the primary adjustment mechanism.

The distinction between jobs and employment is also important when comparing our results to the literature. Notably, the results in BK are often interpreted as saying that migration is the dominant adjustment mechanism, *even in the short run*. This impression is partly based on a misinterpretation of the results in BK. BK do not directly measure population flows across states, but instead attribute the contribution of migration as the share not accounted for by movements in unemployment or participation, as in our decomposition (1). But, instead of using employment data, BK use CES jobs data. The migration component in BK therefore captures not only migration, but also changes in dual job holdings and commuting (as in (6)). When BK use CPS employment data instead of CES jobs data, the migration component falls by 40 percent.

Table 3: LABOR MARKET RESPONSES INCLUDING JOBS

		Short-run ($h = 0$)					Long-run ($h = 5$)				
		Jobs	je	ur	Lfp	Pop	Jobs	je	ur	Lfp	Pop
β	State	0.99 (0.13)	0.30 (0.07)	-0.57 (0.09)	0.03 (0.07)	0.10 (0.04)	2.81 (0.30)	0.21 (0.16)	-0.62 (0.17)	0.39 (0.14)	1.59 (0.23)
	CZ	0.99 (0.05)	0.04 (0.04)	-0.53 (0.04)	0.31 (0.04)	0.11 (0.02)	1.87 (0.18)	0.15 (0.08)	-0.15 (0.07)	0.58 (0.09)	0.99 (0.09)
	County	0.78 (0.03)	0.15 (0.03)	-0.34 (0.02)	0.22 (0.03)	0.07 (0.01)	1.41 (0.08)	0.26 (0.06)	-0.07 (0.03)	0.47 (0.06)	0.62 (0.05)
γ	State		0.31 (0.08)	-0.57 (0.12)	0.03 (0.07)	0.10 (0.05)		0.07 (0.06)	-0.22 (0.06)	0.14 (0.05)	0.57 (0.10)
	CZ		0.04 (0.04)	-0.53 (0.03)	0.31 (0.04)	0.11 (0.02)		0.08 (0.04)	-0.08 (0.03)	0.31 (0.04)	0.53 (0.04)
	County		0.19 (0.04)	-0.43 (0.02)	0.28 (0.03)	0.09 (0.01)		0.19 (0.04)	-0.05 (0.02)	0.33 (0.04)	0.44 (0.03)

Note: The table displays the β 's and γ 's for the Bartik instrument when looking at jobs instead of employment, based on decomposition (6). See also notes to Table 1.

3.1.7 ACCOUNTING FOR OUTSIDE OPTIONS

Our baseline regression ignores that net migration in a given location is driven not only by the labor demand shock to that location but also by the shocks to potential alternative locations (Borusyak et al., 2022a). Even if workers are responsive to local labor demand, there will be little incentive to migrate if workers' current and potential alternative locations face the same labor market conditions. This problem is worse for regions with similar industrial composition and with large migration flows. For instance, consider Ohio and Michigan: two states that both specialize in car manufacturing and have large cross-state migration. A drop in demand for cars will increase unemployment in both states, but there will be little incentive to migrate between the two states. Omitting the outside option from the regression will bias the migration response. In Foschi et al. (2023), we partially address this concern by including a measure of labor demand of a region's "migration partner", calculated as a weighted average across likely destinations and origins of movers. Following this approach, we estimate

$$\Delta \ln POP_{i,t+h} = \alpha_{i,h}^{POP} + \alpha_{t,h}^{POP} + \gamma_h^{POP} \Delta \ln E_{i,t+h} + \gamma_h^{POP,partner} \Delta \ln E_{i,t+h}^{partner} + \Gamma_h^{POP} X_{i,t} + \varepsilon_{i,t+h}^{POP} \quad (7)$$

using as instruments for $\Delta \ln E_{i,t+h}$ and $\Delta \ln E_{i,t+h}^{partner}$ the industry-mix instruments $Z_{i,t}^{IND}$ and $Z_{i,t}^{IND,Partner} \equiv \sum_{p=1}^N \bar{s}_i^p Z_{i,t}^{IND}$, with $\bar{s}_i^p$ indicating the average weight of state p in state i 's set of movers' origins and destinations.




Figure 9: Controlling for labor demand shock in effective partner

Note: The left panel plots the estimated γ_h^{POP} coefficients obtained from running (7) at different horizons h (x-axis), together with the corresponding γ_h^{POP} from the baseline regression that does not control for labor demand shocks in states' effective partner. The right panel plots the estimate of $\gamma_h^{POP,partner}$. The sample is at the state level and runs from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

Figure 9 shows the estimates: Migration responds negatively to employment growth in a state's partner, with a slightly larger absolute effect than its response to local employment growth. However, large standard errors prevent rejecting the null of equal-sized responses. Including partner employment growth in the regression only slightly increases estimates of γ_h^{POP} (e.g. 0.67 vs. 0.61 for $h = 5$). The omitted variable bias in the baseline regression follows an intuitive pattern: Since employment shocks are positively correlated between states and their partners (0.32), the baseline regression understates migration responses. It overlooks that worsening labor market conditions in one state often coincide with weaker conditions in alternative destinations, reducing migration incentives. The magnitude of this bias depends on the volatility of the omitted variable. In our case, after controlling for fixed effects, the standard deviation of a partner state's labor demand shock is only one-sixth that of a state's own shock. This low volatility likely stems from either similar migration patterns across states or weak spatial correlation in industry compositions. Thus, a worker in Michigan facing a labor demand decline will likely see drops elsewhere—but to a lesser extent. While omitting this effect biases migration estimates downward, the bias remains modest in our setup.²⁹

²⁹ Assuming that the true coefficient on the omitted variable has the opposite sign, but is of equal size as the coefficient of interest, then the omitted variable bias in a standard OLS framework is equal to the ratio of the standard deviation of the omitted variable to the standard deviation of the variable of interest times their correlation. In our context, this would yield $0.32 \times \frac{1}{6} \approx 5\%$. Borusyak et al. (2022a) point out that workers are likely to face switching

3.2 ALTERNATIVE INSTRUMENTS FOR LABOR DEMAND SHOCKS

In this section, we consider three alternative proxy variables for labor demand that have received substantial attention in the literature: shocks to regional government spending (i.e. defense contracts) (Nakamura and Steinsson, 2014; Auerbach et al., 2020), shocks to regional housing net worth (Mian and Sufi, 2014; Bhattarai et al., 2021) and regional shocks to import competition (Autor et al., 2013, 2021). We slightly depart from our econometric specification (3) that we used for the industry-mix instrument, mostly to be line with the original papers: Instead of regressing our labor market variables on the labor demand instrument, $Z_{i,t}$, itself, we regress it on the proxy variable for labor demand, which we then instrument for using a shift-share variable. This simplifies the interpretation of the β_h coefficients and leads to the same point estimates of the γ_h coefficients as if we regressed directly on the instrument.³⁰

3.2.1 DEFENSE SPENDING INSTRUMENT

Following Nakamura and Steinsson (2014), we consider regional changes in defense spending as a local labor demand shift variable. We then estimate the response of employment in region i at time t to a change in defense spending with the regression

$$\ln E_{i,t+h} - \ln E_{i,t-2} = \alpha_{i,h}^E + \alpha_{t,h}^E + \beta_h^E \frac{G_{i,t} - G_{i,t-2}}{GDP_{i,t-2}} + \Gamma_h^E X_{i,t} + \varepsilon_{i,t+h}^E. \quad (8)$$

Here, $G_{i,t}$ denotes nominal federal military spending in state i at time t and $GDP_{i,t}$ denotes state i 's nominal GDP. As in Nakamura and Steinsson (2014), we focus on two-year differences to mitigate any timing mismatches between the date of when additional spending is awarded and the effective date of spending.³¹ The coefficient of interest, β_h^E , is the response of employment growth between $t - 2$ and $t + h$ to an increase in government spending by 1% of state GDP between $t - 2$ and t . As we did in the previous section, we also estimate the responses of unemployment, labor force participation, and population to changes in defense spending.

Following Nakamura and Steinsson (2014), we instrument state-level military spending using a shift-share variable. Like the industry Bartik, the military spending instrument is the product

costs across both location and industries. Fully addressing the bias from omitting outside options would therefore require data on workers' transitions between industry-location cells to estimate the outside option. This might lead to a larger estimate of the bias.

³⁰This is true as long as all regressions include the same set of regressors, so that the γ 's add up to 1.

³¹See also Brunet 2022 and Briganti and Sellemi 2023 for a discussion of this timing issue. This approach allows us to partially control for the mismatch between the fiscal year in which the military award data are recorded (starting on October 1st of the previous calendar year) and the calendar year on which the labor market variables are based.

of an aggregate shift term and a regional share term. Specifically, we construct a variable $Z_{i,t}^G$ as the product of the percent change in national military spending (over a two-year horizon) and i 's regional share of government spending to regional GDP. Specifically, for each region i , we calculate the share $s_i = G_i/GDP_i$. The military spending instrument is then,

$$Z_{i,t}^G = \frac{1}{2} \left[\frac{s_{i,t-3}}{s_{US,t-3}} + \frac{s_{i,t-4}}{s_{US,t-4}} \right] \times \frac{G_t - G_{t-2}}{GDP_{t-2}}.$$

Intuitively, states with relatively higher shares of military spending per capita should be affected by changes in national military spending more than states with relative lower shares.

Because variations in the instrument are driven not only by changes in national military spending but also by changes in the exposure shares, we include the shares as separate controls in $X_{i,t}$ (see also Adao et al., 2019; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022b). We also include a lag of the change in federal military spending in state i . Standard errors are clustered at the region level, and regressions are weighted by a region's share in national population.

DATA We consider two regional samples: The first sample is at the state-level (48 states) and covers the years 1976 to 2019. The data on federal military spending by state in Nakamura and Steinsson (2014) covers the years till 2006. We extend their sample to 2019 by collecting military prime contract data from [usaspending.gov](https://www.usaspending.gov).

The second sample follows Auerbach et al. (2020) and is at the level of core-based statistical areas (CBSAs). CBSAs consist of one or more counties that are tied to an urban center by commuting and are therefore conceptually similar to CZs. The data is directly taken from Auerbach et al. (2020) and covers 369 CBSAs for the years 2001 to 2016. Given the short time frame, we calculate the shares for the shift-share instrument using sample averages, as in Auerbach et al. (2020). We therefore do not separately control for these shares because they are absorbed by the region fixed effects.

RESULTS Figure 10 shows the estimated local projections at the state level in response to an increase in defense spending amounting to 1 percent of state GDP. The responses are comparable to those for the industry-mix shock (see Figure 3), in terms of both magnitude and dynamics. In both cases, we observe an employment response that steadily builds up over the first five years. More concretely, an increase in defense spending of 1 percent of a state's GDP raises employment by about 0.55% upon impact, but this response rises to 3% over time. Similarly to the industry-




Figure 10: RESPONSE TO DEFENSE SPENDING INSTRUMENT AT THE STATE LEVEL

Note: The figure plots the beta coefficients in regression 8 for each labor market variable at different horizons h (x-axis) using the CBSA-level dataset (sample period: 1976 - 2019). Shaded areas represent 90% confidence intervals.

mix instrument, roughly 60% of the initial increase in employment is accounted for by reduced state unemployment, while net migration only accounts for 40%. Over time, the migration response becomes stronger and eventually accounts for three-quarters of the adjustment, while the unemployment margin falls to less than 20%.

While the industry-mix instrument has been widely used to proxy for labor demand shifts, it does not easily map into a structural shock: Even though it attenuates the reverse causality issue that a fall in regional population (e.g. due to a positive labor demand shock in a neighboring state, or changing climate conditions) lowers employment—see our discussion above, changes in industries' national employment could still be driven by shifts in regional population patterns, especially if industries are regionally concentrated.³² The defense spending shock is less prone to such criticism. The fact that the two shocks produce very similar responses, though, are suggestive that the industry-mix instruments picks up a fair amount of shifts in labor demand that are uncorrelated with regional population movements.

Figure 11 displays the local projections at the CBSA level. While the initial employment response is similar to the one found at the state level, it does not increase that much over time and

³²The fact that we exclude region i in constructing the Bartik shift mitigates this concern.




Figure 11: RESPONSE TO DEFENSE SPENDING INSTRUMENT AT THE CBSA LEVEL

Note: The figure plots the beta coefficients in regression 8 for each labor market variable at different horizons h (x-axis) using the CBSA-level dataset (sample period: 2001 - 2016). Shaded areas represent 90% confidence intervals.

levels off at 1%. Most of the response is accounted for by a stronger participation rate as well as a lower unemployment rate. In contrast to the baseline and all other labor demand shocks, the population response is negligible and does not build up over time. It is somewhat larger when we focus on working-age population, but even then, it accounts for less than 20% of the employment response.

3.2.2 HOUSING NET WORTH

As a third labor demand shifter, we consider the sudden change in housing net worth during the Great Recession (Mian and Sufi, 2014; Bhattarai et al., 2021). In contrast to the previous two labor demand shifters (industry mix and defense spending) that exploited variation across both regions and time, this third labor demand shifter has only a cross-sectional component.

The main regression to estimate the response of employment to a change in net housing worth




Figure 12: RESPONSE TO HOUSING NET WORTH INSTRUMENT AT THE COUNTY LEVEL

Note: The figure plots the beta coefficients in regression 9 for each labor market variable at different horizons h (x-axis) using the county-level dataset. Horizon $h = 0$ corresponds to year 2007. Shaded areas represent 90% confidence intervals.

follows Bhattacharai et al. (2021)³³

$$\ln E_{i,2007+h} - \ln E_{i,2006} = \alpha_h^E + \beta_h^E \frac{\Delta \ln p_{i,2006-2009}^H \times H_{i,2006}}{NW_{i,2006}} + \sum_{s=1998,2002} \theta_{s,h}^E \bar{g}_{i,s,s-4}^E + \varepsilon_{i,2007+h}, \quad (9)$$

where the housing net worth shock, $\frac{\Delta \ln p_{i,2006-2009}^H \times H_{i,2006}}{NW_{i,2006}}$, consists of the value of housing in region i as of 2006 times the log change in region i 's house prices normalized by a region's overall net worth.

For the previous two labor demand shifters, region fixed effects helped control for any region-specific sample-wide trends in labor market variables (e.g. long-term population shifts across U.S. states). Since we cannot use region fixed effects with cross-sectional data, we follow Bhattacharai et al. (2021) and control for different growth trajectories across regions by including the average pre-trend growth rate over 1994 - 1998 and over 1998 - 2002 of the outcome variable, $\bar{g}_{i,1998,1994}^E$ and

³³As explained by Bhattacharai et al. (2021), there are two differences to Mian and Sufi (2014): First, in Mian and Sufi (2014), the LHS variable is the log difference between 2007 and 2009; second, Mian and Sufi (2014) do not control for pre-trend growth.

$\bar{g}_{i,2002,1998}^E$. Foschi et al. (2023) argue that differences in pre-trends across regions might be particularly important to take into account for net migration as those regions that were particularly hit by the Great Recessions were those that had seen a net inflow of workers in previous years. For our specification here, we note that excluding the pre-trends makes our estimates noisier, but has only modest effects on the point estimates.

We use estimates of the housing supply elasticity from Saiz (2010) as an instrument for the housing net worth shock. As in Bhattarai et al. (2021), we transform the instrument into a dummy that equals 1 if the housing price elasticity is in the upper tercile. This improves the relevance of the instrument.

DATA The data on housing net worth and the Saiz instrument is taken from Mian and Sufi (2014). It only covers larger counties, resulting in $N = 916$.

RESULTS Figure 12 shows the estimated local projections from (9). The impact response, which corresponds to the change in employment between 2006 and 2007 as predicted by housing price changes between 2006 and 2009, is negligible. It takes 2-3 more years for employment dynamics to fully unfold. By 2010, counties that experienced a 10% reduction in housing wealth compared to other counties are predicted to have lost more than 4% of employment more. The response then levels off in subsequent years. The fall in employment goes along with an increase in the unemployment rate and, contrary to our results for all other labor demand shocks, an increase in the labor force participation rate. This labor force response is consistent with a wealth effect on labor supply predicted by neoclassical business cycle models (see also Disney and Gathergood, 2018, for further empirical evidence based on UK data). The migration response is somewhat stronger than for our baseline results. By 2010, close to 60% of the employment change is accounted for by changes in population, with this share steadily increasing to 100% by the middle of the decade.³⁴

3.2.3 IMPORT COMPETITION

Our final labor demand shifter exploits variation in import competition from China across labor markets (Autor et al., 2013). We follow Autor et al. (2021) who consider a purely cross-sectional setup that takes the increase in import penetration between 2000 and 2012 as a labor demand shifter.

³⁴Cadena and Kovak (2016) find that during the Great Recession, low-skilled Mexican workers exhibited a high degree of labor mobility, relocating from the hardest hit regions to less depressed regions.

The main regression to estimate the response of employment to a change in import competition is

$$\ln E_{i,2001+h} - \ln E_{i,2000} = \alpha_h^E + \beta_h^E \sum_j s_{i,2000}^j \frac{IM_{China,2012}^{h,US} - IM_{China,2000}^{j,US}}{Y_{2000}^j + IM_{2000}^j - EX_{2000}^j} + \theta_h^E \bar{g}_{i,1976-1990}^E + \Gamma_h^E X_{i,t} + \varepsilon_{i,2000+h}, \quad (10)$$

where $\frac{IM_{China,2012}^{j,US} - IM_{China,2000}^{j,US}}{Y_{2000}^j + IM_{2000}^j - EX_{2000}^j}$ is the growth of Chinese import penetration for U.S. manufacturing industry j over 2000 - 2012 and $s_{i,2000}^j$ is the share of industry j in region i 's total employment in the base year 2000. The change in import penetration is computed as the growth in U.S. industry imports from China, $IM_{China,2012}^{j,US} - IM_{China,2000}^{j,US}$, divided by initial industry domestic absorption (U.S. industry shipments plus net imports, $Y_{2000}^j + IM_{2000}^j - EX_{2000}^j$).

Similar to the housing-net-worth shock, we control for secular trends in labor market variables across U.S. regions by adding the log change of the left-hand-side variable over 1976-1990 as a regressor. This is particularly important for population: as pointed out by Greenland et al. (2019), CZs that were particularly exposed to import competition from China had had particularly strong population growth in the past. Not accounting for these pre-trends would bias our estimates upwards (towards zero).³⁵

In addition to controlling for the trend growth rate of the dependent variable, we also add the same set of controls as Autor et al. (2021). These include the sum of shares $\sum_j s_{i,1990}^j$ (i.e. the share of employment in manufacturing), initial CZ employment composition, demographic conditions and census region dummies.³⁶ As in Autor et al. (2021), we instrument the measure of import competition using China's exports to other developed countries.

DATA Data on the growth of Chinese import penetration in the United States and the instrument are taken from the replication package of Autor et al. (2021). It covers $N = 719$ CZs.

RESULTS Exposure to import competition from China leads to a decline in employment with an estimated coefficient of around -2.7% in the mid to late 2000s, which then falls to -4% by 2019. Taking into account the interquartile range of 0.56 for trade exposure reported by Autor et al. (2021), this implies that a CZ at the 75th percentile of trade exposure would be predicted to have

³⁵We confirmed this by removing the trend growth rate from our regressions, leading to weaker responses of employment and population.

³⁶Note that we deviate from Autor et al. (2021) by controlling for CZ's share of manufacturing in 1990 rather than 2000 because the instrument uses shares as of 1990. As emphasized by Borusyak et al. (2022b), one needs to control for the sum of shares of the instrument rather than the sum of shares of the endogenous variable. We thank our discussant Brian Kovak for pointing this out.




Figure 13: RESPONSE TO IMPORT COMPETITION INSTRUMENT AT THE CZ LEVEL

Note: The figure plots the beta coefficients in regression 10 for each labor market variable at different horizons h (x-axis) using the CZ-level dataset. Horizon $h = 0$ corresponds to year 2001. Shaded areas represent 90% confidence intervals.

a reduction in employment that is about 2.25 percentage points larger than in a CZ at the 25th percentile. Initially, most of the adjustment occurs via the participation margin, but over time, the population response becomes stronger and accounts for two thirds of the employment response by the mid to late 2000s and more than 80 percent in the long run.³⁷

3.2.4 SUMMARY

We summarize the results for the alternative labor demand shocks in Table 4. The table presents the estimates for the effect of the four labor demand shocks on the labor market variables in the “long run” (i.e., $h = 5$) for each of our instruments. We should acknowledge that the dates at which the import competition shock and the housing price shock begin are somewhat uncertain. As constructed, the import competition shock captures changes in imports that unfold over 13 years (see equation (10)). For the housing-net-worth shock, we take the start date as 2007 and

³⁷In a recent paper, Autor et al. (2025) exploit employer-employee data to decompose changes in employment rates into several changes, including a migration channel. They find a more modest response for the migration channel, but acknowledge that their migration channel only captures workers that were employed in both 2000 and 2019 but in different CZs. Hence, cross-CZ moves of workers that change labor market status or that enter / leave the working-age population would not be captured by their migration channel.

report the response in 2012; for the import competition shock, we take the start date as 2001 and report the response in 2006.

The first part of the table reports the estimated β coefficients that were also depicted in Figures 10 - 13. Since the magnitudes of these labor market responses are not directly comparable across instruments, the γ 's in the bottom of the table are more informative.³⁸ Overall, all the instruments—except for the defense spending instrument at the CBSA level—suggest that migration plays a dominant role in the long run, accounting for about two thirds of the overall employment response. This is consistent with the findings based on the industry composition instrument reported in Table 1.

Table 4: LONG RUN LABOR MARKET RESPONSES FOR ALTERNATIVE LABOR DEMAND INSTRUMENTS

		Empl	ur	Lfp	Pop
β	Defense spending (State)	2.00 (0.71)	−0.50 (0.24)	0.15 (0.37)	1.35 (0.46)
	Defense spending (CBSA)	1.03 (0.59)	−0.44 (0.28)	0.49 (0.48)	0.09 (0.24)
	Housing net worth (County)	0.47 (0.13)	−0.18 (0.04)	−0.09 (0.07)	0.29 (0.05)
	Import competition (CZ)	−2.41 (1.08)	0.03 (0.23)	−1.24 (0.77)	−1.25 (0.59)
γ	Defense spending (State)		−0.25 (0.15)	0.07 (0.19)	0.68 (0.33)
	Defense spending (CBSA)		−0.43 (0.37)	0.48 (0.54)	0.09 (0.24)
	Housing net worth (County)		−0.38 (0.13)	−0.18 (0.15)	0.62 (0.20)
	Import competition (CZ)		−0.01 (0.10)	0.51 (0.40)	0.52 (0.34)

Note: The table displays β_5 and γ_5 for alternative labor demand instruments. For the housing net worth shock and the import competition shock, the responses are measured relative to the year in which the shock is thought to have started. Specifically, for the housing-net-worth shock, the response is measured in 2012 (5 years from 2007); for the import-competition shock, the response is measured in 2006 (5 years from 2001). See also notes to Table 1.

³⁸Note that the γ 's for the housing-net-worth shock and the import-competititon shock do not add up to one because the regressors differ across labor market variables.

4 CONCLUSION

Americans are less likely to relocate to another state now than they were in the past. While some 10 million people – 3.1 percent of the US population - moved to another state in 2021, this is a decline in the cross-state migration rate relative to earlier decades. One interpretation of this decline is that the U.S. job market is less dynamic than it used to be. Our findings call this interpretation into question. We estimate the responsiveness of unemployment, labor force participation and labor migration to exogenous variations in labor demand. We also estimate the labor migration elasticity, the fraction of the total employment adjustment attributed to changes in employment.

Our baseline estimates are based on the now-standard industry-mix instrument (Bartik, 1993). Echoing the earlier findings of BK, we find that cross-state labor mobility accounts for most of the labor adjustment to the shock at five-year horizons. While migration became less responsive to labor demand shocks in the 1990s (Dao et al., 2017), so did employment. The elasticity of migration to employment fell little and has recovered since. Estimates of the migration elasticity are similar based on data from commuting zones and counties. Not surprisingly, the migration elasticity is largest for individuals of working age, though it does not differ by gender or race.

We confirm that labor mobility was and is a key adjustment mechanism in response to labor demand using alternative instruments: a fiscal shock, the Great Recession shock and the China shock. At five-year horizon, cross-region labor mobility accounts for most of the adjustment in employment changes for all shocks except for fiscal shocks when focusing at the CBSA level.

Although the migration elasticity is robust across shocks and in most time periods, we do not want to overstate our results. Most workers remain in place, even in bad times. But our estimates suggest that the elasticity of migration to shocks has remained fairly constant over time, and is as strong as it was in the 1950s. This is not to discount the many studies that suggest that the labor market now is less flexible than in the past. Movements in and out of the labor market, shifts between jobs, and the underlying correlation of shocks to the labor market – over time and across regions – have likely changed. What our findings do suggest is that geographic relocation remains an important channel of labor market adjustment.

REFERENCES

- R. Adao, M. Kolesár, and E. Morales. Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010, 2019.
- M. Amior and A. Manning. The persistence of local joblessness. *American Economic Review*, 108(7):1942–1970, 2018.
- J. Angrist, P. Hull, and C. Walters. Methods for measuring school effectiveness. *Handbook of the Economics of Education*, 7:1–60, 2023.
- J. D. Angrist and P. Hull. Instrumental variables methods reconcile intention-to-screen effects across pragmatic cancer screening trials. *Proceedings of the National Academy of Sciences*, 120(51):e2311556120, 2023.
- A. J. Auerbach, Y. Gorodnichenko, and D. Murphy. Macroeconomic frameworks: Reconciling evidence and model predictions from demand shocks. *NBER Working Paper*, 26365, 2020.
- D. Autor, D. Dorn, and G. H. Hanson. On the persistence of the china shock. *Brookings papers on economic activity*, 2021(3):381–447, 2021.
- D. Autor, D. Dorn, G. H. Hanson, M. R. Jones, and B. Setzler. Places versus people: the ins and outs of labor market adjustment to globalization. Technical report, National Bureau of Economic Research, 2025.
- D. H. Autor, D. Dorn, and G. H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, 103(6):2121–2168, 2013.
- T. J. Bartik. Who benefits from local job growth: migrants or the original residents? *Regional studies*, 27(4):297–311, 1993.
- P. Beaudry, D. A. Green, and B. M. Sand. The declining fortunes of the young since 2000. *American Economic Review*, 104(5):381–386, 2014.
- M. Beraja, E. Hurst, and J. Ospina. The aggregate implications of regional business cycles. *Econometrica*, 87(6):1789–1833, 2019.
- S. Bhattarai, F. Schwartzman, and C. Yang. Local scars of the us housing crisis. *Journal of Monetary Economics*, 119:40–57, 2021.
- O. J. Blanchard and L. F. Katz. Regional evolutions. *Brookings papers on economic activity*, 1992(1):1–75, 1992.
- K. Borusyak, R. Dix-Carneiro, and B. Kovak. Understanding migration responses to local shocks. *Available at SSRN 4086847*, 2022a.
- K. Borusyak, P. Hull, and X. Jaravel. Quasi-experimental shift-share research designs. *The Review of economic studies*, 89(1):181–213, 2022b.
- K. Borusyak, P. Hull, and X. Jaravel. A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204, 2025.
- J. Bound and H. J. Holzer. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of labor Economics*, 18(1):20–54, 2000.
- E. Briganti and V. Sellemi. Why does gdp move before government spending? it’s all in the measurement. *UCSD Manuscript*, 2023.
- G. Brunet. When does government spending matter? evidence from a new measure of spending. *en. In: Working Paper (May 2020)*, page 41, 2022.

- B. C. Cadena and B. K. Kovak. Immigrants equilibrate local labor markets: Evidence from the great recession. *American Economic Journal: Applied Economics*, 8(1):257–290, 2016.
- M. Dao, D. Furceri, and P. Loungani. Regional labor market adjustment in the united states: Trend and cycle. *Review of Economics and Statistics*, 99(2):243–257, 2017.
- J. DeWaard, M. Hauer, E. Fussell, K. J. Curtis, S. D. Whitaker, K. McConnell, K. Price, D. Egan-Robertson, M. Soto, and C. Anampa Castro. The emergence of a uniform business cycle in the united states: Evidence from new claims-based unemployment data. *Population Research and Policy Review*, 41(2):437–448, 2022.
- R. Disney and J. Gathergood. House prices, wealth effects and labour supply. *Economica*, 85(339):449–478, 2018.
- F. Eckert, T. C. Fort, P. K. Schott, and N. J. Yang. Imputing missing values in the us census bureau’s county business patterns. Technical report, National Bureau of Economic Research, 2020.
- A. J. Fieldhouse, S. Howard, C. Koch, and D. Munro. The emergence of a uniform business cycle in the united states: Evidence from new claims-based unemployment data. *Brookings Papers on Economic Activity*, forthcoming, 2024.
- A. Foote, M. Grosz, and A. Stevens. Locate your nearest exit: Mass layoffs and local labor market response. *ILR Review*, 72(1):101–126, 2019.
- A. Foschi, C. L. House, C. Proebsting, and L. L. Tesar. Labor mobility and unemployment over the business cycle. In *AEA Papers and Proceedings*, volume 113, pages 590–596. American Economic Association, 2023.
- P. Goldsmith-Pinkham, I. Sorkin, and H. Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.
- A. Greenland, J. Lopresti, and P. McHenry. Import competition and internal migration. *Review of Economics and Statistics*, 101(1):44–59, 2019.
- N. Jia, R. Molloy, C. Smith, and A. Wozniak. The economics of internal migration: Advances and policy questions. *Journal of economic literature*, 61(1):144–180, 2023.
- G. Kaplan and S. Schulhofer-Wohl. Understanding the long-run decline in interstate migration. *International Economic Review*, 58(1):57–94, 2017.
- A. Mian and A. Sufi. What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223, 2014.
- R. Molloy, R. Trezzi, C. L. Smith, and A. Wozniak. Understanding declining fluidity in the us labor market. *Brookings Papers on Economic Activity*, 2016(1):183–259, 2016.
- R. Molloy, C. L. Smith, and A. Wozniak. Job changing and the decline in long-distance migration in the united states. *Demography*, 54(2):631–653, 2017.
- J. Monras. Economic shocks and internal migration. 2018.
- E. Nakamura and J. Steinsson. Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *The American Economic Review*, 104(3):753–792, 2014.
- M. J. Notowidigdo. The incidence of local labor demand shocks. *Journal of Labor Economics*, 38(3):687–725, 2020.
- D. G. O’Connor. Revitalize or relocate: Optimal place based transfers for local recessions. 2024.

- J. R. Pierce and P. K. Schott. Trade liberalization and mortality: evidence from us counties. *American Economic Review: Insights*, 2(1):47–63, 2020.
- A. Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.
- D. Sullivan and T. Von Wachter. Job displacement and mortality: An analysis using administrative data. *The Quarterly Journal of Economics*, 124(3):1265–1306, 2009.
- R. Wilson. Isolated states of america: State borders, mobility, and labor markets. *Employment Research Newsletter*, 29(1):2, 2022.
- D. Yagan. Moving to opportunity? migratory insurance over the great recession. 2014.