ANDREA FOSCHI
Bank of Italy

CHRISTIAN PROEBSTING
KU Leuven

CHRISTOPHER L. HOUSE
University of Michigan

LINDA L. TESAR
University of Michigan

Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks

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 by a reduction in unemployment. Over time however, net labor in-migration becomes the dominant factor contributing to increased regional employment. After five years, roughly 60 percent of the increase in employment is explained by the change in population. Responses of labor migration are strongest for individuals age 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.

Conflict of Interest Disclosure: Andrea Foschi's employer, Bank of Italy, had the right to review this paper before publication but did not influence the findings. The authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. The authors are not currently an officer, director, or board member of any organization with a financial or political interest in this paper. The Brookings Institution is committed to quality, independence, and impact. It is supported by a diverse array of funders. In line with its values and policies, each Brookings publication represents the sole views of the author(s).

Brookings Papers on Economic Activity, Spring 2025: 67–114 © 2025 The Brookings Institution.

In 2010, the population of Williston, North Dakota, was 14,716 people—up from 12,512 a decade earlier, according to the Census. 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 United States—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 640,000. The ebb and flow of workers across regions is clearly part of the dynamism of the US labor market. In a well-known study of labor market dynamics in 1992, Blanchard and Katz (1992) 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 Blanchard and Katz (1992), 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 US states has declined over the last forty 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 US labor market flexibility (Molloy and others 2016).

The fact that migration rates have declined, however, does not necessarily mean that labor mobility has become less responsive to regional

1. Jia and others (2023) provide a comprehensive overview of recent research on changes in US migration patterns.

Figure 1. Gross Cross-State Migration Rate, 1976–2021

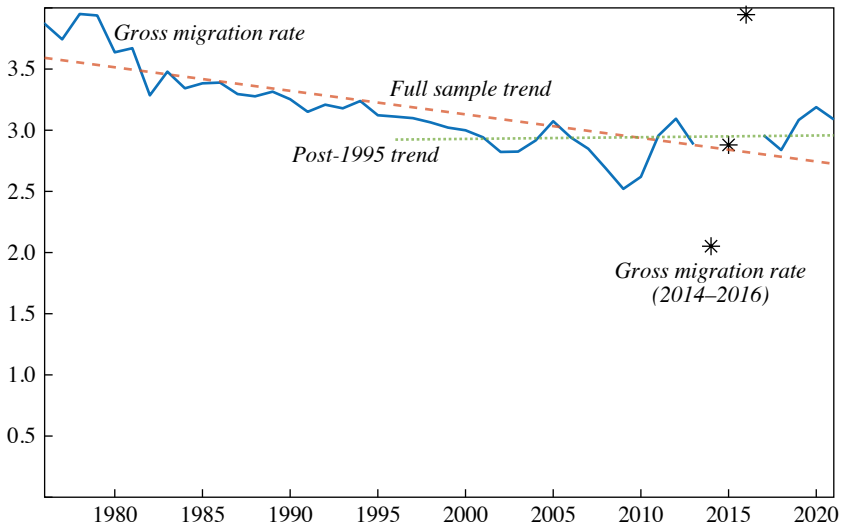

Source: IRS and authors' calculations.

Note: The figure plots the gross cross-state migration rate, based on IRS data, computed as $\frac{1}{48}\sum_{i=1}^{48} 0.5 \frac{In\text{-}migrants_{i,t} + Out\text{-}migrants_{i,t}}{Non\text{-}migrants_{i,t} + Out\text{-}migrants_{i,t}}$, that is, 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 and others 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 grayed out in the plot. Nonetheless, we believe post-2012 data are still informative to have a sense of overall trends.

economic disturbances. In this paper we examine the joint responsiveness of labor migration, labor force participation, and unemployment to plausibly 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 mix instrument, based on Bartik (1993) and similar to Dao, Furceri, and Loungani (2017); a military spending share instrument, based on Nakamura and Steinsson (2014) and Auerbach, Gorodnichenko, and Murphy (2024); an import share instrument, that is, the "China shock," based on Autor, Dorn, and Hanson (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 data set 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 five-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 body of literature on the extent of labor mobility in the United States. While Blanchard and Katz (1992) argue that labor mobility was the dominant adjustment mechanism, subsequent studies find a smaller role for regional labor mobility (Dao,

Furceri, and Loungani 2017; Fieldhouse and others 2024). Studies focused on the Great Recession (e.g., Mian and Sufi 2014; Yagan 2014) and on the effects of Chinese import penetration (e.g., Autor, Dorn, and Hanson 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, Hurst, and Ospina 2019).

Not all recent research supports the view that labor migration has faded in importance, however. For instance, Foote, Grosz, and Stevens (2019, p. 102) study the effects of mass layoffs on regional employment and find that "the predominant adjustment mechanism following a mass layoff event is migration." However, like other researchers, Foote, Grosz, and Stevens (2019) find that migration played a substantially smaller role during and after the Great Recession. In recent work, Foschi and others (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]

I. 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

2. In recent work, O'Connor (2025) studies the intertemporal trade-off between policies that support workers in hard-hit areas and providing incentives to relocate in the long run.

3. Note 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.

in the unemployment rate, labor force participation rate, and regional population (i.e., migration). After discussing the data for these variables, we present ordinary least squares (OLS) estimates on the historical comovement 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.

I.A. Decomposition of Employment Growth

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

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

where $U_{i,t}$ denotes the total number of people unemployed, $LF_{i,t} = E_{i,t} + U_{i,t}$ denotes the labor force, and $POP_{i,t}$ denotes the population.[4] 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]

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

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

4. Employment 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 II.A.

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

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]

I.B. 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.

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 are 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 ($N = 48$).

6. The 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 $1 + nm_{i,t}$. Some studies have linked the number of births and deaths to local business cycles: For instance, Black and others (2013) show that fertility increased in the Appalachian coal mining regions following the 1970s increase in world energy prices; Schaller, Fishback, and Marquardt (2020) document a positive association between local per capita income growth and the subsequent growth of local birth rates; while Sullivan and Von Wachter (2009) and Pierce and Schott (2020) report somewhat higher mortality among workers facing earnings losses. Pierce and Schott (2020) in particular report that counties more exposed to import competition from China (one of our shocks) see a slight increase in "deaths of despair," often associated with drug overdoses. Quantitatively, however, these effects are small, with an interquartile shift in counties' exposure being associated with a relative increase in mortality by one 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.

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

LABOR MARKET DATA For each county, we obtain data on employment, unemployment, and labor force participation 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 are available at the county level starting in 1976.[9] Annual figures are computed as averages of the January-to-December monthly figures.

In extensions of our baseline analysis, we integrate this employment measure with two additional measures: the employment figures from the Bureau of Economic Analysis (BEA) and from the State and Metro Area Employment, Hours, and Earnings Statistics, contained in the BLS's Current Employment Statistics (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 of 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.

POPULATION AND MIGRATION DATA For each county, we take population data from the BEA.[10] These are available at the county level since 1969 and at the state level since 1929. Population figures are reported as of

8. Employment 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 US 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.

9. Official 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 affected by these changes.

10. This corresponds to the population data produced by the Census Bureau, consisting of decennial counts for census years and intercensal estimates for noncensus 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

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 National Cancer Institute and has the same time and geographic coverage as the BEA population data.[11] Second, the IRS has data on tax return addresses, which is one of the inputs for the production of postcensal estimates, and results in a separate data set 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 into immigration and emigration and to calculate gross migration rates.[13]

I.C. 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

---

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 IRS, and the ACS. The IRS tax return data also form the basis for the IRS migration flow estimates described later.

11. See National Bureau of Economic Research, "Survey of Epidemiology and End Results (SEER) U.S. State and County Population Data by Age, Race, Sex, Hispanic 1969-On," https://www.nber.org/research/data/survey-epidemiology-and-end-results-seer-us-state-and-county-population-data-age-race-sex-hispanic.

12. Estimates 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 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 are 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, and we find that the population response is the same.

13. The ACS provides another source of information about labor migration between states. As we noted in Foschi and others (2023), problems with the collection of ACS data make them a less reliable data source for the study of migration than the IRS/Census data.

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:

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

for different horizons $h = 0, 1, 2, \ldots$ and years t. The i regions are US 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 a given variable $Y_{i,t}$):

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

Note that the hat applies to the whole log-differenced variable, for example, $Y_{i,t} = \Delta \ln POP_{i,t}$. 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 variations. We do this because many regions have persistently high (or low) unemployment rates and persistent in- or out-migration.

Figure 2, panel A, shows the estimated coefficient $\beta_{h,t}$ for $h = 0$ for each year over the 1945–2021 period. Data on employment by state are only available since 1976. To expand the analysis backward, we also consider regressing population changes on changes in the number of jobs, for which

Figure 2. Coefficient of Net Migration on Jobs and Employment Growth, 1946–2021

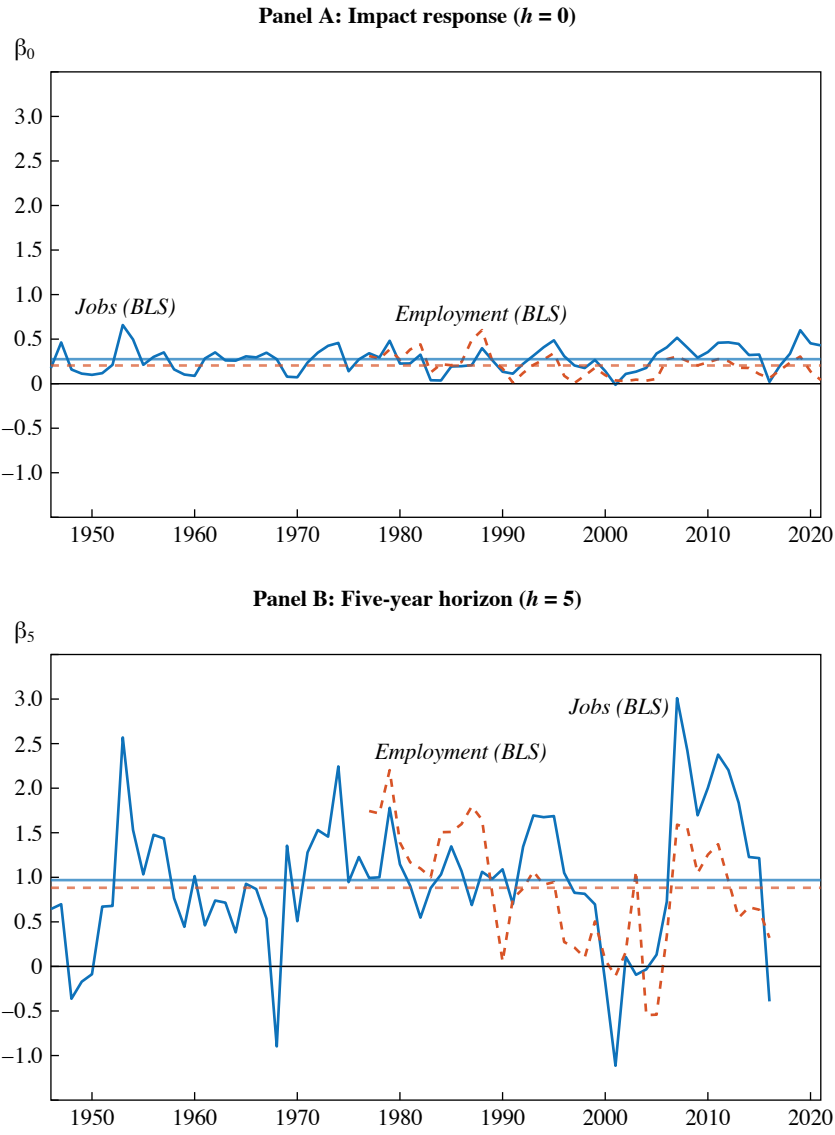

Source: Authors' calculations.

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

longer time series are available. Figure 2, panel B, shows the same year-by-year coefficients for the five-year $h = 5$ cumulative population change. The average contemporaneous coefficient in panel A 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 1 percent above average tend to have current net in-migration equal to roughly 0.2–0.25 percent of the population. The migration effect clearly becomes stronger over longer horizons (panel B). At the $h = 5$ year horizon, the estimated coefficients are approximately one.

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.

I.D. 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 section II.B.

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 $h = 0, 1, \ldots, H$, and for any variable on the left-hand side $Y_{i,t}$, we estimate the following regression:

$$Y_{i,t+h} = \alpha_{i,h}^{Y} + \alpha_{t,h}^{Y} + \beta_{h}^{Y} Z_{i,t} + \Gamma_{h}^{Y} X_{i,t} + \varepsilon_{i,t+h}^{Y}. \tag{3}$$

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 variable $Y_{i,t}$ other than employment itself, we calculate:[14]

$$\gamma_h^Y = \frac{\beta_h^Y}{\beta_h^E}. \tag{4}$$

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 3 percent increase in regional employment was achieved through a 2 percent increase in net migration and a 1 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 calculation 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 one for each horizon h regardless of the econometric specification and provided that the variables on the right-hand side in equation (3) are the same for each independent variable Y:

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

14. We use the delta method to calculate the corresponding standard errors. To obtain the necessary estimates of the variance-covariance matrix, we stack the data set, saturate it with a sample indicator, and run all regressions together (Angrist, Hull, and Walters 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 two-stage least squares (2SLS) approach are always smaller.

15. Generally, 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 percent and its unemployment rate is 5 percent. 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 percent of the adjustment to the population margin.

Of course, the β coefficients and the γ coefficients are simply the first-stage and second-stage estimates from a two-stage least squares (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 United States 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.

II. Estimates of the Migration Elasticity

In this section we present our estimates of the migration elasticity, γ^{POP}. Section II.A presents results for the traditional Bartik instrument, which takes shifts in industrial composition as the regional labor demand instrument. We discuss the results in this section at some length before moving on to the other instruments. Section II.B considers other common instruments for labor demand.

II.A. Industry Mix Instrument

We begin by considering the industry mix instrument originally proposed by Bartik (1993) and later used by Dao, Furceri, and Loungani (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:

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

In this equation, $\varepsilon_{t,-i}^{j}$ 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, that is, $X_{i,t}$ 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 others, Bound and Holzer 2000; Beaudry, Green, and Sand 2014; Notowidigdo 2020). Recent literature in applied microeconometrics has helped clarify under what conditions these shift-share instruments yield consistent estimates (Adão, Kolesár, and Morales 2019; Goldsmith-Pinkham, Sorkin, and Swift 2020; Borusyak, Hull, and Jaravel 2022). Most of the literature focuses on the single cross-section case, which raises concerns that instruments might be correlated with alternative explanations. The panel 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, Hull, and Jaravel 2025).

Borusyak, Hull, and Jaravel (2025) point out that exogeneity of the time series variation in the shifts is sufficient for consistent estimation. Since

16. We also consider alternative ways of constructing the shares: We first average them over the previous nine years (so that, for example, the shares for 1990 are constructed as the average of the 1981–1989 shares); we then average them over the first five years of the previous nine years, skipping the last four years (so that the shares for 1990 are constructed as the average of the 1981–1985 shares). The results are robust.

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

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 the 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 II.B, we therefore look at alternative labor demand instruments and generally find similar results.

Data. We construct the Bartik instrument 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 (CBP), in the version provided by Eckert and others (2020), to construct the shares.[18] As in the data from the BEA and in 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 instrument is the much richer level of industry detail that they allow: In creating the Bartik instrument, we use data at the highest available level of industry detail, which means six digits when using the North American Industry Classification System (NAICS) and four digits when using the Standard Industrial Classification (SIC) system; 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

18. While 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 data from Eckert and others (2020) to compute employment shares at the county level. A key advantage of the Eckert and others (2020) data set is in 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 and others (2020) version of the CBP data give a more complete picture of the detailed industry structure of counties.

approximately one thousand industries overall.[19] We build CZ- and state-level versions of the Bartik instrument as employment-weighted averages of the counties they contain.

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 panel D. 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 90 percent confidence intervals. 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,

19. For the shift, using the QCEW, we rely on SIC data until 1997, the last year when they were available; starting from 1998, we use NAICS data. To ensure there is no unexpected jump between 1997 and 1998, when the industry classification changes, we compute the 1997–1998 employment change using reconstructed 1997 NAICS data, also provided by the BLS. Similarly, for the shares, using CBP data, we rely on the concordance tables constructed by Eckert and others (2020) to convert all industry codes to the newest classification whenever we are averaging across years that include a classification change. For instance, in 1997, the shift variable is based on SIC, and we accordingly compute the average shares over 1993–1997 using the native SIC shares used in those years; however, in 1998, the shift variable switches to NAICS, and we thus 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.

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

Figure 3. Response to Bartik Instrument at the State Level

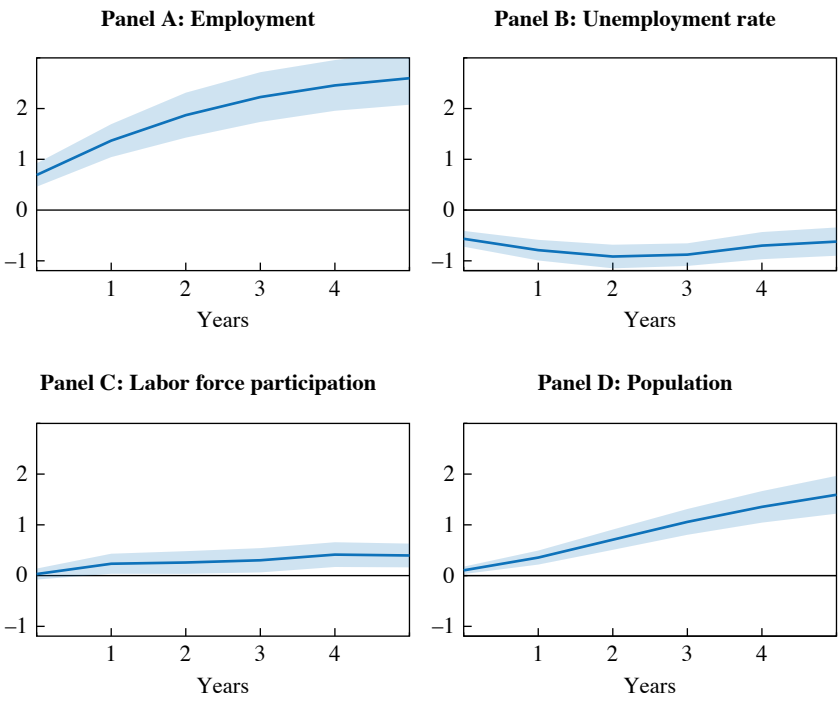

Source: Authors' calculations.

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

these shares switch. At the five-year horizon, more than 60 percent of the estimated increase in employment is due to increased state population while only 20–25 percent comes from reduced unemployment.

It is worth comparing the results to the literature. Based on a vector autoregression (VAR) identified by the assumption that current shocks to employment growth are driven by labor demand only, Blanchard and Katz (1992) 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, Furceri, and Loungani (2017), the identifying assumption in Blanchard and Katz (1992) 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

Figure 4. Ratio Between Each Labor Market Response and the Employment Response

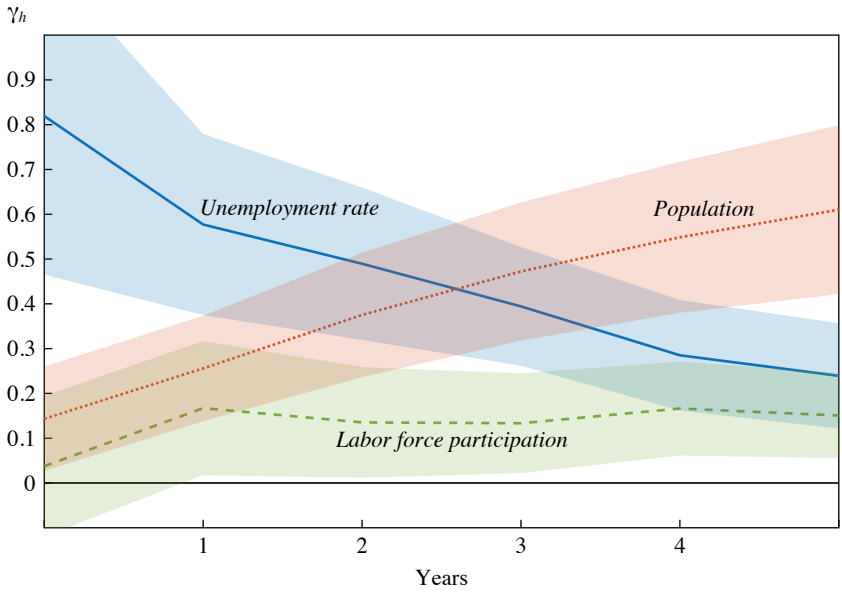

Source: Authors' calculations.

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

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 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, Furceri, and Loungani (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 versus 11 percent; see their table 1), in line with our γ estimates in figure 4.

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 on shorter subsamples 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:

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

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 instrument: From 1976 onward, it corresponds to our baseline four- and six-digit Bartik instrument constructed from the QCEW and the CBP; pre-1976, it corresponds to a one-digit Bartik instrument, 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 ten-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.

21. A few issues prevent us from exactly replicating the baseline decomposition from figures 3 and 4 with pre-1976 historical data: The four- and six-digit Bartik instrument is only available starting 1976; the data on employment (based on number of employees by residence) are only available starting in 1976 (hence why we use jobs here); and the data on unemployment are similarly only available starting in 1976. Consequently, labor force participation (which requires the number of employed plus the number of unemployed) is also only available starting in 1976. As a result, population is the only variable from figures 3 and 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 and others 2024), but labor force participation would still not be available.

Panels A and B of figure 5 show 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 5, panel C.[22] Most of the variation in the γ^{POP} coefficient arises during periods when employment responses were historically low and estimated with less precision, leading to noisier estimates of the γ^{POP}.[23]

This analysis places recent research on labor mobility into a historical context. For example, comparing the original Blanchard and Katz (1992) sample (1976–1990) with subsequent data, Dao, Furceri, and Loungani (2017, p. 251) conclude that the "migration sensitivity to regional shocks in the short and long run, has been strongly decreasing since the 1990s." This is consistent with the β^{POP} coefficient depicted in figure 5, panel B: for the 1976–1990 samples (midpoints 1980.5–1985.5), the β^{POP} coefficient remains consistently above two, 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 figure 5, panel C, is substantially smaller—and by the 2000s it has fully recovered.[24] Viewed from the perspective of

22. The 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 because we considered the response of the number of jobs in the denominator rather than the employment response. See below for a discussion on this distinction between jobs and employment.

23. We check whether there is any cyclicality in γ^{POP} by inspecting its lead-lag correlations with GDP growth and by regressing it on an NBER recession dummy; in both cases there seems to be no clear evidence that γ^{POP} is either procyclical or countercyclical.

24. Note that this result does not contradict figure 1: It depicts gross migration, which reflects unconditional mobility and includes a lot of offsetting flows with no net effect on population; figure 5 instead focuses on economically motivated flows that induce population changes. The findings may be reconciled by a decline in noneconomic migration motives or a decline in the incidence of local labor demand shocks.

Figure 5. Historical Jobs and Population Responses to Bartik Instrument at the State Level, Ten-Year Rolling Windows

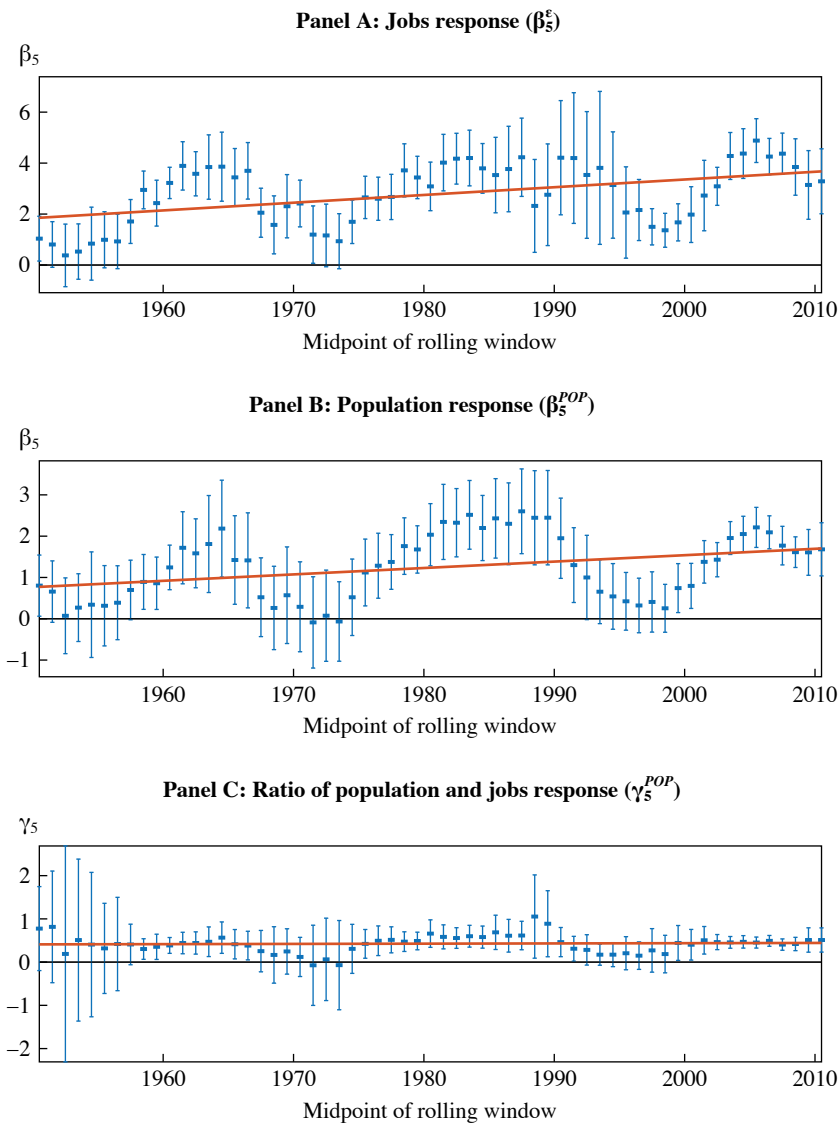

Source: Authors' calculations.

Note: The figure plots the estimates of the β_h coefficients obtained by running equation (5) at horizon $h = 5$ using double-demeaned state-level data over ten-year rolling windows (full sample period: 1946–2016); the γ_h coefficient defined in equation (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 percent confidence intervals. The solid straight line is a weighted least square fit, where the weights are the inverse of the standard errors.

history, there is therefore little evidence that the contribution of migration to labor adjustment has diminished.[25]

MIGRATION RESPONSES FOR SMALLER REGIONS: CZs 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 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 five-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.[26] Nevertheless, for all regional aggregates, net migration becomes the dominant adjustment 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.

25. It also follows that the sum of the contributions of the remaining adjustment margins—labor force participation rate, unemployment, and the job-to-employment margin—has remained reasonably constant over time, but with some variation. This implies that the response of these margins to labor demand shocks has somewhat increased over time, similar to the increased responsiveness of jobs and population depicted in figure 5, panels A and B.

26. These 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 county-level data are suppressed for confidentiality reasons (Eckert and others 2020). Note that this measurement error cancels out when calculating γ. We also find that labor demand changes are less persistent at a more granular level, which would also contribute to a smaller response.

Table 1. Labor Market Responses at Different Levels of Aggregation

		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)

Source: Authors' calculations.

Note: The top half of the table displays the estimates of the β_h coefficients obtained by running equation (3) for each of the dependent variables listed at the top, with $h = 0$ in the left columns and $h = 5$ in the right columns. 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 to 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.

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.[27]

Although the differences are not large, the point estimates suggest that migration elasticities are somewhat smaller at finer geographic levels (see also Foschi and others 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.[28] Using IRS data, we separate county-level migration flows into cross-state and within-state net migration flows and rerun 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.

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 most likely to relocate.[29] 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).

27. These elasticities at the five-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 CZ 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.

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

29. Kaplan and Schulhofer-Wohl (2017) and Molloy, Smith, and Wozniak (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.

Figure 6. Population Response by Age

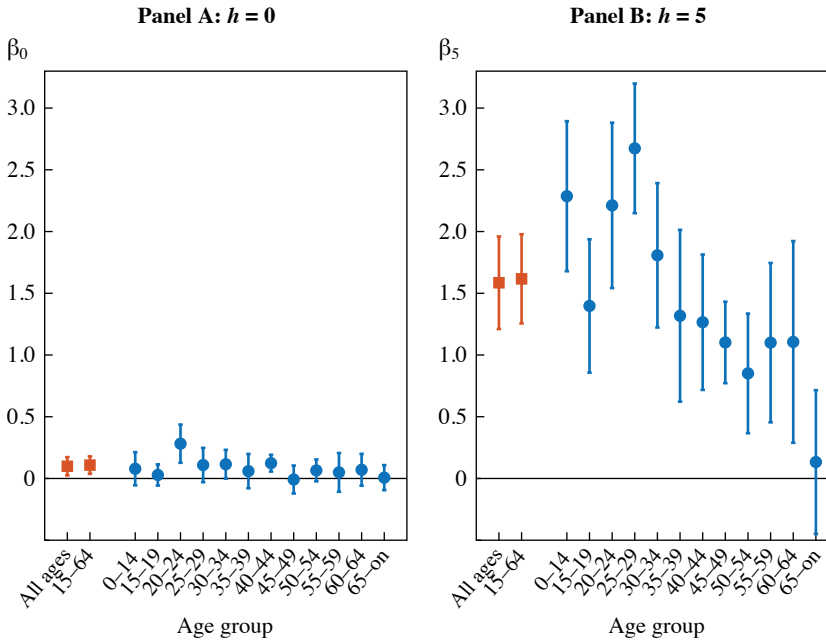

Source: Authors' calculations.

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

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.[30] 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, nonmetropolitan but adjacent to a metropolitan county, or nonmetropolitan with an urban population, according to the categorization provided in the 2013 Rural-Urban

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

Continuum Codes (RUCC) by the US Department of Agriculture (USDA). We code a county as higher income if its income per 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 counties. For urban versus 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 counties 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 counties and 0.49 for less educated counties).

ASYMMETRIES IN MIGRATION RESPONSES We modify our baseline regression to examine two types of asymmetries: (1) differences in the response to positive versus negative shocks and (2) differences in in-migration versus out-migration. As in the previous section, we conduct the analysis at the state level.

To distinguish between positive and negative shocks, we introduce an indicator variable into our estimation equation that differentiates between the two, and we allow the β coefficient to vary based on the sign of the shock.[31] 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.

31. Specifically, we estimate $Y_{i,t+h} = \alpha^{Y}_{i,h} + \alpha^{Y}_{t,h} + \beta^{Y,+}_{h} Z_{i,t} \mathbb{1}_{\{Z_{i,t} \geq 0\}} + \beta^{Y,-}_{h} Z_{i,t} \mathbb{1}_{\{Z_{i,t} < 0\}} + \Gamma^{Y}_{h} X_{i,t} + \varepsilon^{Y}_{i,t+h}$.

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	0.65 (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)

Source: Authors' calculations.

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

Figure 7. Labor Market Responses to Positive and Negative Labor Demand Shocks

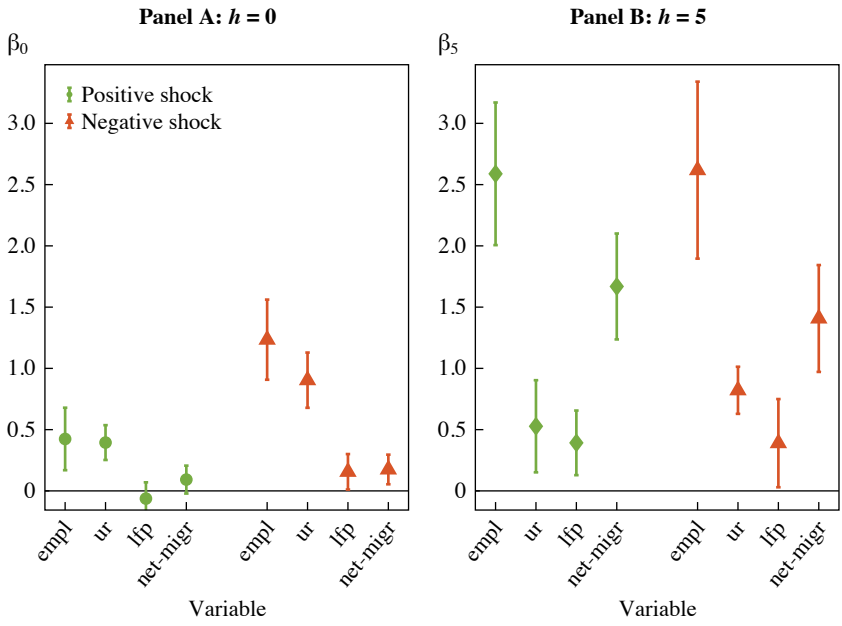

Source: Authors' calculations.

Note: The figure plots the estimated β_h coefficients obtained from running equation (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 to 2016, with projections going up to 2021 (for $h = 5$). The bars indicate 90 percent confidence intervals.

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 (2020), 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.[32]

32. Monras (2020) 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

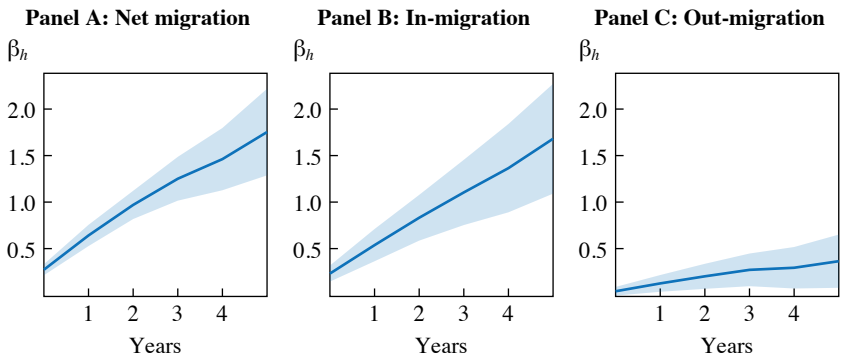

Source: Authors' calculations.

Note: The figure plots the estimated β_h coefficients obtained from running equation (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 to 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90 percent confidence intervals.

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:

$$(6) \qquad \Delta \ln \varepsilon_{i,t} = \Delta \ln \frac{\varepsilon_{i,t}}{E_{i,t}} + \Delta \ln \left(1 - ur_{i,t}\right) + \Delta \ln LFP_{i,t} + \Delta \ln POP_{i,t},$$

where $\varepsilon_{i,t}$ denotes jobs in region i at time t. Table 3 presents the β and γ estimates from reestimating our local projections (3) at the state, CZ, and county level but now including estimates for β_h^{ε} and γ_h^{ε}.

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 (je) accounts for nearly one-third of short-run labor market adjustment at the state level. While the data do not distinguish

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)

Source: Authors' calculations.

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.

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 Blanchard and Katz (1992) 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 their results. Blanchard and Katz (1992) 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, Blanchard and Katz (1992) use CES jobs data. The migration component in their findings therefore captures not only migration, but also changes in dual job holdings and commuting, as in our equation (6). When Blanchard and Katz (1992) use CPS employment data instead of CES jobs data, the migration component falls by 40 percent.

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, Dix-Carneiro, and Kovak 2022). 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 and others (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:

$$\Delta \ln POP_{i,t+h} = \alpha_{i,h}^{POP} + \alpha_{t,h}^{POP} + \gamma_{h}^{POP} \Delta \ln E_{i,t+h} + \gamma_{h}^{POP,partner} \Delta \ln E_{i,t+h}^{partner} + \Gamma_{h}^{POP} X_{i,t} + \varepsilon_{i,t+h}^{POP}, \tag{7}$$

Figure 9. Controlling for Labor Demand Shock in Effective Partner

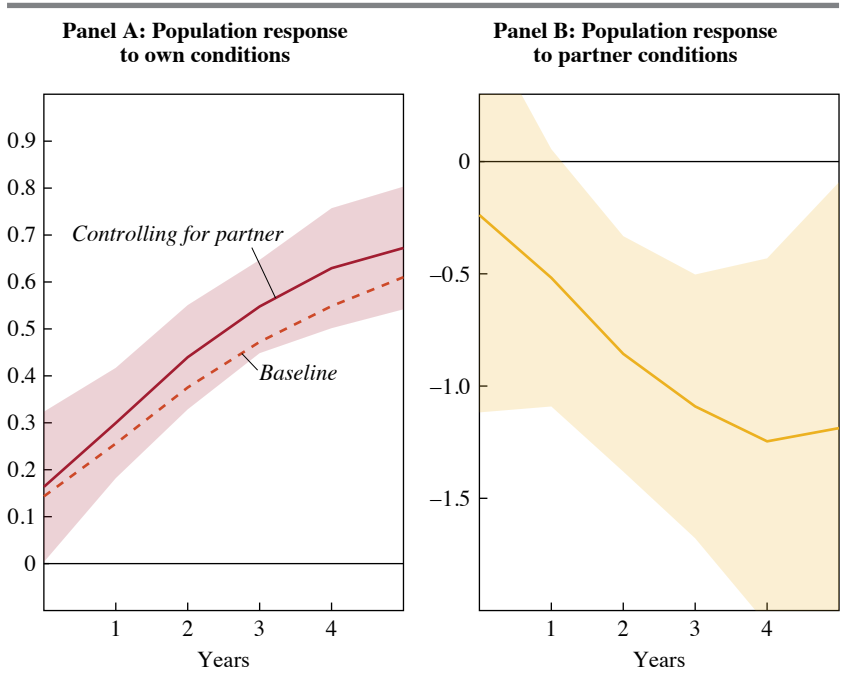

Source: Authors' calculations.

Note: Panel A plots the estimated γ_h^{POP} coefficients obtained from running equation (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. Panel B plots the estimate of $\gamma_h^{POP,partner}$. The sample is at the state level and runs from 1976 to 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90 percent confidence intervals.

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.[33]

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

33. The weights $\bar{s}_i^p$ are computed as the average of two other sets of weights, $\bar{s}_i^{p,IN}$ and $\bar{s}_i^{p,OUT}$, which in turn represent the importance for i of p as a migration origin and a migration destination, respectively. Specifically, $\bar{s}_i^{p,IN}$ is the average share of migrants into i that come from p, and $\bar{s}_i^{p,OUT}$ is the average share of migrants out of i that go to p, computed from bilateral migration flows from the ACS. The results are similar if we replace $\bar{s}_i^p$ with $\bar{s}_i^{p,IN}$ only or with $\bar{s}_i^{p,OUT}$ only.

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 versus 0.61 for $h = 5$).[34] 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.[35]

II.B. 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, that is, defense contracts (Nakamura and Steinsson 2014; Auerbach, Gorodnichenko, and Murphy 2024); shocks to regional housing net worth (Mian and Sufi 2014; Bhattarai, Schwartzman, and Yang 2021); and regional shocks to import competition (Autor, Dorn, and Hanson 2013, 2021). We slightly depart from our econometric specification (3) that we used for the industry mix instrument, mostly to be in 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

34. We also rerun equation (7) at the CZ and county level and find that the γ^{POP}, after controlling for partner conditions, is similarly close to the baseline.

35. 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 $0.32 \times 1/6 \approx 5\%$. Borusyak, Dix-Carneiro, and Kovak (2022) point out that workers are likely to face switching costs across both location and industry. 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.

to the same point estimates of the γ_h coefficients as if we regressed directly on the instrument.[36]

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:

$$\ln E_{i,t+h} - \ln E_{i,t-2} = \alpha^E_{i,h} + \alpha^E_{t,h} + \beta^E_h \frac{G_{i,t} - G_{i,t-2}}{GDP_{i,t-2}} + \Gamma^E_h X_{i,t} + \varepsilon^E_{i,t+h}. \tag{8}$$

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.[37] The coefficient of interest β^E_h is the response of employment growth between $t-2$ and $t+h$ to an increase in government spending by 1 percent 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 instrument, the military spending instrument is the product of an aggregate shift term and a regional share term. Specifically, we construct a variable $Z^G_{i,t}$ as the product of the percentage 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:

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

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

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

37. See 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 1 of the previous calendar year) and the calendar year on which the labor market variables are based.

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 Adão, Kolesár, and Morales 2019; Goldsmith-Pinkham, Sorkin, and Swift 2020; Borusyak, Hull, and Jaravel 2022). 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 (forty-eight states) and covers the years 1976 to 2019. The data on federal military spending by state in Nakamura and Steinsson (2014) cover the years until 2006. We extend their sample to 2019 by collecting military prime contract data from USASpending.gov.

The second sample follows Auerbach, Gorodnichenko, and Murphy (2024) 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 are directly taken from Auerbach, Gorodnichenko, and Murphy (2024) and cover 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, Gorodnichenko, and Murphy (2024). 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 percent upon impact, but this response rises to 3 percent over time. Similarly to the industry mix instrument, roughly 60 percent of the initial increase in employment is accounted for by reduced state unemployment, while net migration only accounts for 40 percent. 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 percent.

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, changes in industries' national employment could still be driven by shifts in regional population

Figure 10. Response to Defense Spending Instrument at the State Level

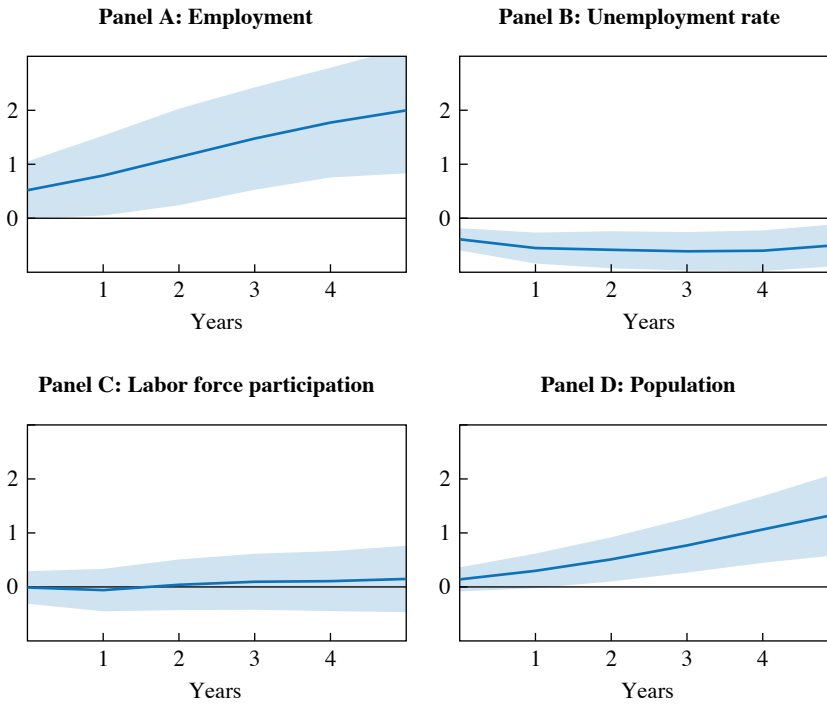

Source: Authors' calculations.

Note: The figure plots the β coefficients in regression (8) for each labor market variable at different horizons h (x-axis) using the state-level data set (sample period: 1976–2019). Shaded areas represent 90 percent confidence intervals.

patterns, especially if industries are regionally concentrated.[38] The defense spending shock is less prone to such criticism. The fact that the two shocks produce very similar responses, though, is suggestive that the industry mix instruments pick 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 levels off at 1 percent. Most

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

Figure 11. Response to Defense Spending Instrument at the CBSA Level

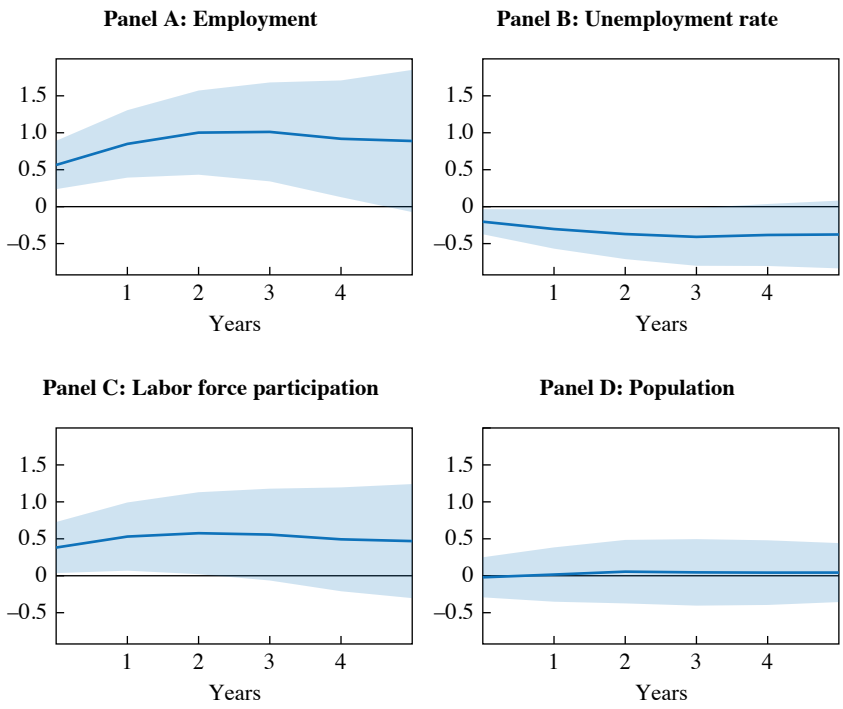

Source: Authors' calculations.

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

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 percent of the employment response.

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, Schwartzman, and Yang 2021). In contrast to the previous two labor demand shifters (industry mix and defense spending) that exploited variation across both region 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 follows Bhattarai, Schwartzman, and Yang (2021):[39]

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

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 US states). Since we cannot use region fixed effects with cross-sectional data, we follow Bhattarai, Schwartzman, and Yang (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 variables, $\bar{g}_{i,1998,1994}^E$ and $\bar{g}_{i,2002,1998}^E$. Foschi and others (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 Recession 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, Schwartzman, and Yang (2021), we transform the instrument into a dummy that equals one 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 are taken from Mian and Sufi (2014). It only covers larger counties, resulting in $N = 916$.

39. As explained by Bhattarai, Schwartzman, and Yang (2021), there are two differences from Mian and Sufi (2014): First, in Mian and Sufi (2014), the variable on the left-hand side is the log difference between 2007 and 2009; second, Mian and Sufi (2014) do not control for pre-trend growth.

Figure 12. Response to Housing Net Worth Instrument at the County Level

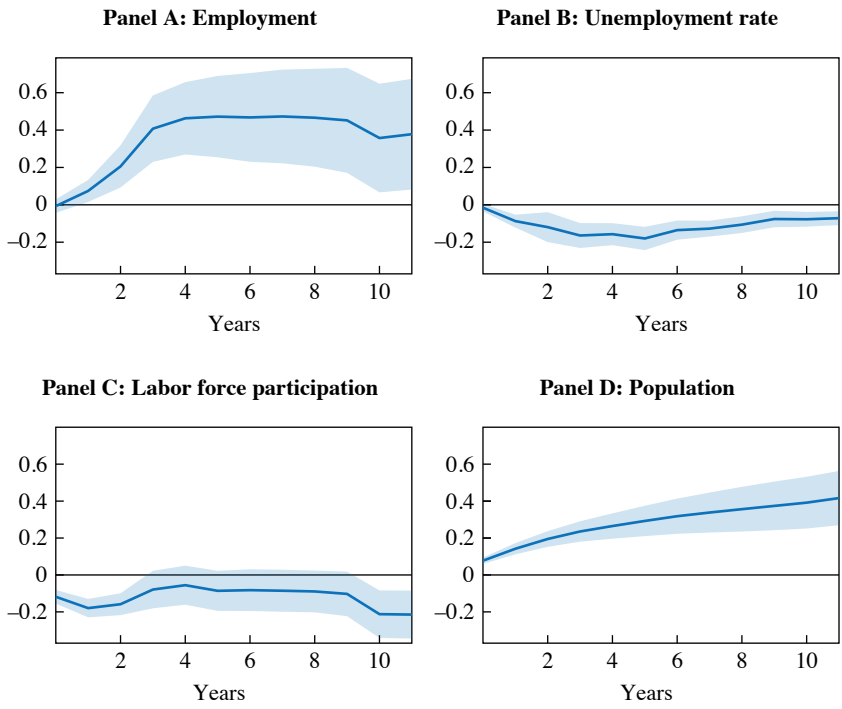

Source: Authors' calculations.

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

Results. Figure 12 shows the estimated local projections from equation (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 two to three more years for employment dynamics to fully unfold. By 2010, counties that experienced a 10 percent reduction in housing wealth compared to other counties are predicted to have lost more than 4 percent of employment. 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 percent of the employment change is accounted for by changes in population, with this share steadily increasing to 100 percent by the middle of the decade.[40]

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

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

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

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 US 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 US industry imports from China, $IM_{China,2012}^{j,US} - IM_{China,2000}^{j,US}$, divided by initial industry domestic absorption (US 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 US regions by adding the log change of the variable on the left-hand side over 1976–1990 as a regressor. This is particularly important for population: As pointed out by Greenland, Lopresti, and McHenry (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 upward (toward zero).[41]

In addition to controlling for the trend growth rate of the dependent variable, we also add the same set of controls as Autor, Dorn, and Hanson (2021).

40. Cadena 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.

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

These include the sum of shares $\sum_j s^j_{i,1990}$ (i.e., the share of employment in manufacturing), initial CZ employment composition, demographic conditions, and census region dummies.[42] As in Autor, Dorn, and Hanson (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, Dorn, and Hanson (2021). It covers $N = 719$ CZs.

Results. The results are shown in figure 13. Exposure to import competition from China leads to a decline in employment with an estimated coefficient of around −2.7 percent in the mid- to late 2000s, which then falls to −4 percent by 2019 (note the plots have a flipped sign for comparability to the other shocks). Taking into account the interquartile range of 0.56 for trade exposure reported by Autor, Dorn, and Hanson (2021), this implies that a CZ at the 75th percentile of trade exposure would be predicted to have 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.[43]

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 thirteen years; see equation (10). For the housing net worth shock, we take the start date as 2007 and report the response in 2012; for the import competition shock, we take the start date as 2001 and report the response in 2006.

42. Note that we deviate from Autor, Dorn, and Hanson (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, Hull, and Jaravel (2022), 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.

43. In a recent paper, Autor and others (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 who change labor market status or who enter or leave the working age population would not be captured by their migration channel.

Figure 13. Response to Import Competition Instrument at the CZ Level

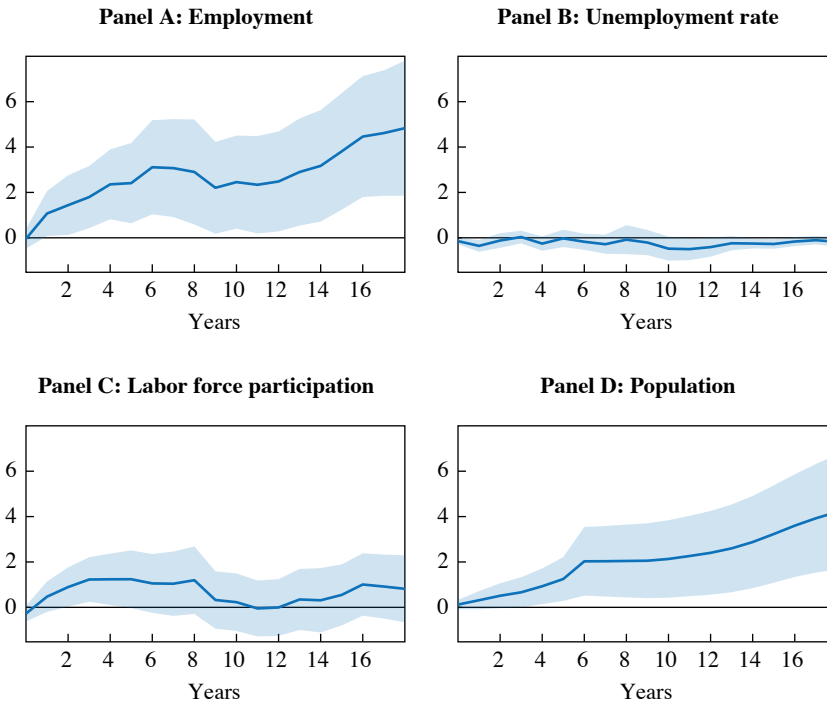

Source: Authors' calculations.

Note: The figure plots the β coefficients in regression (10) for each labor market variable at different horizons h (x-axis) using the CZ-level data set. Horizon $h = 0$ corresponds to year 2001. Shaded areas represent 90 percent confidence intervals. Relative to the other shocks, the China shock has an opposite sign, so for comparability we flip the responses in this plot.

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.[44] 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.

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

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)	0.89 (0.59)	–0.38 (0.28)	0.47 (0.47)	0.04 (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)		–0.42 (0.42)	0.53 (0.63)	0.05 (0.28)
	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)

Source: Authors' calculations.

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 (five years from 2007); for the import competition shock, the response is measured in 2006 (five years from 2001). See also notes to table 1.

III. Conclusion

Americans are less likely to relocate to another state now than they were in the past. According to our own calculations based on the IRS data, while about nine million people—2.8 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 US 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 Blanchard and Katz (1992), we find that cross-state labor mobility accounts for most of the labor adjustment to the shock at five-year horizon. While migration became less responsive to labor demand shocks in the 1990s (Dao, Furceri, and Loungani 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 CZs 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 the 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.

ACKNOWLEDGMENTS We gratefully acknowledge financial support from the Michigan Institute for Teaching and Research in Economics (MITRE). We also thank Brian Kovak, Abigail Wozniak, and Jón Steinsson for their comments and suggestions. The views expressed here are those of the authors and do not necessarily reflect those of the Bank of Italy or the Eurosystem.

References

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-Share Designs: Theory and Inference." *Quarterly Journal of Economics* 134, no. 4: 1949–2010.

Amior, Michael, and Alan Manning. 2018. "The Persistence of Local Joblessness." *American Economic Review* 108, no. 7: 1942–70.

Angrist, Joshua D., and Peter Hull. 2023. "Instrumental Variables Methods Reconcile Intention-to-Screen Effects Across Pragmatic Cancer Screening Trials." *Proceedings of the National Academy of Sciences* 120, no. 51: e2311556120.

Angrist, Joshua D., Peter Hull, and Christopher Walters. 2023. "Methods for Measuring School Effectiveness." *Handbook of the Economics of Education* 7: 1–60.

Auerbach, Alan J., Yuriy Gorodnichenko, and Daniel Murphy. 2024. "Macroeconomic Frameworks: Reconciling Evidence and Model Predictions from Demand Shocks." *American Economic Journal: Macroeconomics* 16, no. 3: 190–229.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103, no. 6: 2121–68.

Autor, David, David Dorn, and Gordon H. Hanson. 2021. "On the Persistence of the China Shock." *Brookings Papers on Economic Activity*, Fall: 381–447.

Autor, David, David Dorn, Gordon H. Hanson, Maggie R. Jones, and Bradley Setzler. 2025. "Places Versus People: The Ins and Outs of Labor Market Adjustment to Globalization." Working Paper 33424. Cambridge, Mass.: National Bureau of Economic Research.

Bartik, Timothy J. 1993. "Who Benefits from Local Job Growth: Migrants or the Original Residents?" *Regional Studies* 27, no. 4: 297–311.

Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2014. "The Declining Fortunes of the Young Since 2000." *American Economic Review* 104, no. 5: 381–86.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2019. "The Aggregate Implications of Regional Business Cycles." *Econometrica* 87, no. 6: 1789–833.

Bhattarai, Saroj, Felipe Schwartzman, and Choongryul Yang. 2021. "Local Scars of the US Housing Crisis." *Journal of Monetary Economics* 119: 40–57.

Black, Dan A., Natalia Kolesnikova, Seth G. Sanders, and Lowell J. Taylor. 2013. "Are Children 'Normal'?" *Review of Economics and Statistics* 95, no. 1: 21–33.

Blanchard, Olivier J., and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, no. 1: 1–75.

Borusyak, Kirill, Rafael Dix-Carneiro, and Brian K. Kovak. 2022. "Understanding Migration Responses to Local Shocks." Working Paper. Social Science Research Network, April 14. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4086847.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies* 89, no. 1: 181–213.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. "A Practical Guide to Shift-Share Instruments." *Journal of Economic Perspectives* 39, no. 1: 181–204.

Bound, John, and Harry J. Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s." *Journal of Labor Economics* 18, no. 1: 20–54.

Briganti, Edoardo, and Victor Sellemi. 2023. "Why Does GDP Move Before Government Spending? It's All in the Measurement." Working Paper, May 23. https://edoardobriganti.github.io/research/paper_draft_MAR2023_whydoesgdp.pdf.

Brunet, Gillian. 2022. "When Does Government Spending Matter? Evidence from a New Measure of Spending." Working Paper, May 31. https://economics.ucdavis.edu/sites/g/files/dgvnsk13091/files/media/documents/6-2%20Brunet.pdf.

Cadena, Brian C., and Brian K. Kovak. 2016. "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession." *American Economic Journal: Applied Economics* 8, no. 1: 257–90.

Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. "Regional Labor Market Adjustment in the United States: Trend and Cycle." *Review of Economics and Statistics* 99, no. 2: 243–57.

DeWaard, Jack, Mathew Hauer, Elizabeth Fussell, Katherine J. Curtis, Stephan D. Whitaker, Kathryn McConnell, Kobie Price, David Egan-Robertson, Michael Soto, and Catalina Anampa Castro. 2022. "User Beware: Concerning Findings from the Post 2011–2012 U.S. Internal Revenue Service Migration Data." *Population Research and Policy Review* 41: 437–48.

Disney, Richard, and John Gathergood. 2018. "House Prices, Wealth Effects and Labour Supply." *Economica* 85, no. 339: 449–78.

Eckert, Fabian, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang. 2020. "Imputing Missing Values in the US Census Bureau's County Business Patterns." Working Paper 26632. Cambridge, Mass.: National Bureau of Economic Research.

Fieldhouse, Andrew J., David Munro, Christoffer Koch, and Sean Howard. 2024. "The Emergence of a Uniform Business Cycle in the United States: Evidence from New Claims-Based Unemployment Data." *Brookings Papers on Economic Activity*, Spring: 265–319.

Foote, Andrew, Michel Grosz, and Ann Stevens. 2019. "Locate Your Nearest Exit: Mass Layoffs and Local Labor Market Response." *ILR Review* 72, no. 1: 101–26.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar. 2023. "Labor Mobility and Unemployment over the Business Cycle." *AEA Papers and Proceedings* 113: 590–96.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110, no. 8: 2586–624.

Greenland, Andrew, John Lopresti, and Peter McHenry. 2019. "Import Competition and Internal Migration." *Review of Economics and Statistics* 101, no. 1: 44–59.

Jia, Ning, Raven Molloy, Christopher Smith, and Abigail Wozniak. 2023. "The Economics of Internal Migration: Advances and Policy Questions." *Journal of Economic Literature* 61, no. 1: 144–80.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration." *International Economic Review* 58, no. 1: 57–94.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82, no. 6: 2197–223.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the US Labor Market." *Brookings Papers on Economic Activity*, Spring: 183–237.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25, no. 3: 173–96.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2017. "Job Changing and the Decline in Long-Distance Migration in the United States." *Demography* 54, no. 2: 631–53.

Monras, Joan. 2020. "Economic Shocks and Internal Migration." Discussion Paper DP12977. London: Centre for Economic Policy Research.

Nakamura, Emi, and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104, no. 3: 753–92.

Notowidigdo, Matthew J. 2020. "The Incidence of Local Labor Demand Shocks." *Journal of Labor Economics* 38, no. 3: 687–725.

O'Connor, Daniel G. 2025. "Revitalize or Relocate: Optimal Place-Based Transfers for Local Recessions." Working Paper, January 14. https://oconnordanny.github.io/files/RevitalizeRelocate.pdf.

Pierce, Justin R., and Peter K. Schott. 2020. "Trade Liberalization and Mortality: Evidence from US Counties." *American Economic Review: Insights* 2, no. 1: 47–64.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125, no. 3: 1253–96.

Schaller, Jessamyn, Price Fishback, and Kelli Marquardt. 2020. "Local Economic Conditions and Fertility from the Great Depression Through the Great Recession." *AEA Papers and Proceedings* 110: 236–40.

Sullivan, Daniel, and Till Von Wachter. 2009. "Job Displacement and Mortality: An Analysis Using Administrative Data." *Quarterly Journal of Economics* 124, no. 3: 1265–306.

Wilson, Riley. 2022. "Isolated States of America: State Borders, Mobility, and Labor Markets." *Employment Research* 29, no. 1: 4–7.

Yagan, Danny. 2014. "Moving to Opportunity? Migratory Insurance over the Great Recession." Working Paper, January. https://eml.berkeley.edu/~yagan/MigratoryInsurance.pdf.

Comments and Discussion

COMMENT BY

BRIAN KOVAK In a large and diverse modern economy, the mix of occupations and industries varies substantially across locations, due to both natural cost advantages and agglomeration spillovers (Ellison and Glaeser 1999). In this context, economic shocks affecting particular occupations or sectors take on a spatial dimension, and regional markets may face explicitly geographic shocks such as natural disasters or local policy changes. In both cases, local economic shocks drive differences in labor market outcomes in different locations, incentivizing workers to migrate toward local markets facing the most favorable shocks. By migrating, these workers improve their own economic outcomes, equilibrate spatial differences in labor market conditions for nonmigrants, and improve economic efficiency (Bartik 1991; Blanchard and Katz 1992; Borjas 2001; Cadena and Kovak 2016).

Yet, there is increasing concern that these potential economic benefits are being lost due to secular declines in interregional migration in the United States. A large body of empirical literature documents consistent reductions in migration probabilities since at least the 1980s (e.g., Molloy, Smith, and Wozniak 2011; Kaplan and Schulhofer-Wohl 2012). This reduction in the migration rate has occurred simultaneously with reductions in job creation, job destruction, and job-to-job transitions (Davis, Faberman, and Haltiwanger 2012; Hyatt and Spletzer 2013), raising concerns about a broader decline in dynamism in the US labor market (Molloy and others 2016). However, a decline in unconditional migration rates does not necessarily imply a reduction in migration responsiveness to economic shocks, and direct evidence on the latter is mixed. In their seminal paper covering

Brookings Papers on Economic Activity, Spring 2025: 115–137 © 2025 The Brookings Institution.

1976–1990, Blanchard and Katz (1992) find substantial population responses to state-level shocks, with local employment rates equilibrating within six years. Dao, Furceri, and Loungani (2017) update and extend the Blanchard and Katz results through 2013 and find steady reductions in migration responsiveness from the early 1990s onward. In contrast, Yagan (2019) finds that, although local employment rates took much longer to converge after the Great Recession than in the recessions of the early 1980s and 1990s, population responses were similar in the two time periods and in line with Blanchard and Katz (1992). Because these papers use different data sources, empirical models, and sources of identifying variation, it is difficult to generate strong conclusions about how migration responsiveness has changed over time.

Foschi, House, Proebsting, and Tesar fill this gap in the literature by studying US migration responses to economic shocks across seven decades, using comparable data and a consistent set of empirical approaches. Their main analysis examines the effects of a standard Bartik (1991) industry shift-share shock on regional employment, population, labor force participation, and unemployment using a variation on a stacked difference-in-differences specification. The authors document effects at different time horizons relative to the shock and show how the effects for a fixed relative time horizon evolve over calendar time. While the paper presents the effects on each of the four outcomes, its main focus is the "migration elasticity," defined as the shock's effect on population divided by its effect on employment.

The paper generates two headline findings. First, immediately following a local labor demand shock, employment changes are primarily accounted for by changes in local unemployment, while in the longer run migration accounts for a majority of the employment effect. Second, the long-run (five-year) migration elasticity remained relatively constant from 1946 to 2016. Based on these results, the authors argue that migration remains a centrally important spatial adjustment mechanism and that its influence shows no signs of waning over time.

This paper makes an important contribution by providing seventy years of consistent and comparable evidence on population and reemployment responses to labor demand shocks. Simply assembling the necessary data to implement such a long-term analysis is an impressive feat. The paper fills in and expands the historical record on local responses to shocks while avoiding the difficulties in comparing across papers that use different data sources and empirical methods. The authors also avoid the potential pitfalls associated with specific local labor demand shock measures by using four

different measures drawn from prior work. In addition to the Bartik-style industry shift share, they use alternative shock measures based on local military spending, import competition, and changes in housing values. While only the Bartik shock has sufficient time coverage to support arguments about changes in responsiveness across decades, the alternative shock measures confirm the finding that migration accounts for a majority of estimated employment effects at the five-year time horizon. These features make this paper essential reading for anyone studying regional labor markets.

My comments will focus on three aspects of the paper and its relationship to the broader literature. First, I will argue that the migration elasticity may be less informative regarding declining dynamism than the direct effect of shocks on migration. Second, I demonstrate the practical importance of pre-trends in identifying migration effects and show how the paper's approach to addressing pre-trends differs from standard specifications in the broader literature. Third, I will discuss some potential explanations for the paper's surprising finding that estimated effects of shocks on population growth are smaller when using more detailed region definitions.

MIGRATION ELASTICITY INTERPRETATION A novel element of this paper is its focus on the migration elasticity, the effect of a local demand shock on the local population divided by the effect on local employment. This measure is attractive for two main reasons. First, it places the effects of different shock measures in the same units, those of employment growth, facilitating comparison across alternative identification strategies. Second, the proportional change in local employment can be additively decomposed into proportional changes in population, labor force participation, and the unemployment rate. Due to the properties of linear regression, the estimated effect of a shock on employment can also be additively decomposed into effects on the other three elements, providing a helpful accounting of the dimensions through which employment changes occurred. It is important to note, however, that this decomposition does not have a causal interpretation—it is purely an accounting decomposition. Finding that the effect on population growth accounts for 50 percent of the effect on employment growth does not imply that the effect on employment growth would have been 50 percent smaller if migration was not possible.

How, then, should one interpret estimates of the migration elasticity, and what do they teach us about declining dynamism? To answer these questions, consider a simplified version of the local labor markets model of Amior and Manning (2018), who develop a Rosen-Roback-style model with an equilibrium housing sector, frictional local population adjustment,

and endogenous labor supply among those residing in a given location. I simplify their dynamic population adjustment process by replacing it with an isoelastic relationship between local population and local real wages, with elasticity η. Consider the effect of an exogenous vector of shocks to local labor demand, dz_r^d, across regions r. Imposing market clearing equilibrium in both housing and labor markets yields the following value for the migration elasticity, γ^{Pop}.[1]

$$\gamma^{Pop} \equiv \frac{d\ell_r / dz_r^d}{dn_r / dz_r^d} = \frac{\eta}{\eta + \epsilon^{ns}},$$

where ℓ_r is log population, n_r is log employment, and ϵ^{ns} is the labor supply elasticity among those residing in r. This expression highlights two important points. First, an increase in the population elasticity, η, increases both the numerator and denominator of γ^{Pop}, since a larger population response to shocks drives a larger overall employment response. Second, the migration elasticity depends on both the population elasticity η and the labor supply elasticity ϵ^{ns} conditional on population, so changes over time in γ^{Pop} reflect changes in both underlying elasticities. If both elasticities fall, say because an aging population reduces both η and ϵ^{ns}, the migration elasticity γ^{Pop} may be roughly constant over time despite the clear decline in economic dynamism. As shown in figure 5 in the paper, a situation along these lines occurred in the 1990s, leading prior authors such as Dao, Furceri, and Loungani (2017) to conclude that migration responses fell during this period.

Given this interpretation, the direct population response to shocks is more informative regarding migration's role in changing economic dynamism than is the migration elasticity. Happily, the authors present the employment and population responses in figure 5, so readers can see how the numerator and denominator of γ^{Pop} evolve separately. There are large

1. Following the notation of Amior and Manning (2018), replace their dynamic migration model with a static isoelastic local population supply curve $\ell_r = \eta(w_r - p_r)$. Equating labor demand and supply in Amior and Manning's equations (5) and (6), differentiating with respect to the local labor demand shock z_r^d, and combining with the local population supply curve yields:

$$\frac{d\ell_r}{dz_r^d} = \frac{\eta}{\epsilon^{ns} + \eta}\left(1 - \epsilon^{nd}\frac{dw_r}{dz_r^d}\right).$$

Differentiating the labor demand curve shows that the term in parentheses equals dn_r/dz_r^d, yielding the result for the migration elasticity.

Figure 1. Effect of the China Shock on Population Growth, Age 25–39, Alternative Approaches to Pre-Trends

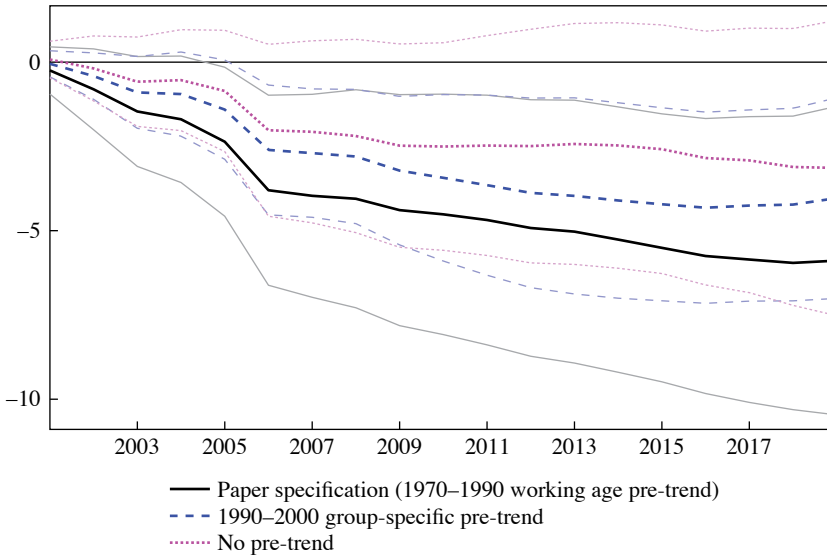

Source: Author's calculations using replication files and data from Autor, Dorn, and Hanson (2021).
Note: Showing alternative approaches to addressing correlated preshock trends in population growth. Thinner lines show 95 percent confidence intervals associated with the estimates shown in bold lines with the same pattern.

differences in estimated population responses from 1950 to 2010, with estimates ranging from zero to greater than 2.5. Understanding the drivers of these large differences in migration responsiveness across the decades should be a priority for future work.

PRE-TRENDS A pervasive issue in the literature studying the effects of local labor demand shocks on migration is the presence of correlated pre-trends in population growth. The decision of how to address pre-trends can have important consequences for the quantitative and qualitative findings regarding the effects of local shocks on contemporaneous population growth. As an example, consider the population analysis in Autor, Dorn, and Hanson (2021), who measure the effects of increased imports from China on population growth across US commuting zones from the year 2000 onward. The authors replicate and expand this analysis in section II.B of their paper and note that controlling for pre-trends substantially affects their findings.

To see the quantitative impact of different approaches to controlling for pre-trends in this context, figure 1 here replicates figure 7, panel C,

in Autor, Dorn, and Hanson (2021), which examines commuting zone population growth for those age 25–39. I focus on this sample because it is the only age group for which Autor, Dorn, and Hanson consistently find statistically significant population effects. The solid line replicates the analysis of Autor, Dorn, and Hanson, who address pre-trends by controlling for the change in log working age (18–64) population from 1970 to 1990. Note that this pre-trend control does not include information in the decade 1990–2000, which just precedes the outcome window, and applies to the overall working age population rather than the specific population under study, those age 25–39. The dashed line omits this older pre-trend control and replaces it with the 1990–2000 change in log population for workers age 25–39. The effect magnitudes in this specification are smaller by roughly one-third. Finally, the dotted line omits the pre-trend controls all together, yielding effects that are nearly 50 percent smaller in magnitude than the baseline specification and no longer statistically significantly different from zero at the 5 percent level. While I do not claim that any one of these approaches is superior to the others, the substantial quantitative differences across specifications make clear the importance of appropriately addressing the presence of pre-trends in this context.

How, then, do the authors deal with the issue of pre-trends in their main Bartik shift-share analysis? Their equation (5) relates the double-demeaned change in log population to the double-demeaned Bartik shock (see section I.C for details). As the authors note, this double demeaning is akin to including location and time fixed effects, as in the following specification (omitting the control vector X for expositional clarity).

$$\Delta_{t-1}^{t+h} \ln Pop_i = \alpha_h^{Pop} + \beta_h^{Pop} Z_{i,t} + \eta_i + \varphi_t + \varepsilon_{i,t+h}^{Pop},$$

where Δ_a^b represents the change from time a to b, η_i are location fixed effects, and φ_t are time fixed effects; all other notations follow the paper. To interpret this specification, difference across adjacent time periods within location:

$$\Delta_{t-1}^{t+h} \ln Pop_i - \Delta_{t-2}^{t-1+h} \ln Pop_i = \left(\varphi_t - \varphi_{t-1}\right) + \beta_h^{Pop}\left(Z_{i,t} - Z_{i,t-1}\right) + \left(\varepsilon_{i,t+h}^{Pop} - \varepsilon_{i,t-1+h}^{Pop}\right).$$

This formulation makes clear the counterfactual assumption underlying the paper's main specification: Local population growth would persist (plus an additive term common across locations) when the Bartik shock remains constant (e.g., when national industry growth is constant). Finally, to relate

the paper's specification to the most common approach in the prior literature, move the population pre-trend, $\Delta_{t-2}^{t-1+h} \ln Pop_i$, to the right-hand side:

$$\Delta_{t-1}^{t+h} \ln Pop_i = \left(\varphi_t - \varphi_{t-1}\right) + \beta_h^{Pop}\left(Z_{i,t} - Z_{i,t-1}\right) + 1 \cdot \Delta_{t-2}^{t-1+h} \ln Pop_i + \left(\varepsilon_{i,t+h}^{Pop} - \varepsilon_{i,t-1+h}^{Pop}\right).$$

This specification differs from the standard approach in the literature along two dimensions. First, it relates local population growth not to the Bartik shock facing the region, but to the *change* in that shock. Second, it restricts the coefficient on the pre-trend to be one, rather than allowing it to take on an arbitrary value, which would allow for imperfect persistence in population growth absent shocks.

Both of these changes may substantially affect the estimated relationship between shocks and population growth, and they make it difficult to compare the results to prior work, which tends to find much smaller effects of local shocks on population growth. Understanding the implications of these specification differences should be a priority for future work. Moreover, given the ubiquity of correlated pre-trends in this literature, future work will benefit from more structured approaches to understanding the sensitivity of results to pre-trend violations, such as the method developed by Rambachan and Roth (2023).

SMALLER EFFECTS IN SMALLER REGIONS When studying the effects of local labor demand shocks, it is necessary to choose a level of geographic aggregation. Most recent papers studying the United States define regions as commuting zones, which are defined as markets in which the vast majority of individuals live and work in the same zone. The main analysis of Bartik shocks in the paper uses state-level geography, which is much more aggregate, but the authors helpfully present results on the employment decomposition using commuting zone and county-level geography. For all three levels of aggregation, they find that changes in unemployment account for the majority of employment growth in the year of the shock but that migration accounts for the majority five years later.[2]

While the relative shares are consistent across levels of aggregation, the magnitudes of the effects on each element of the decomposition are generally smaller when using more detailed geography. As the authors emphasize, this pattern is surprising for migration, since one would expect moves between smaller locations to be less costly and therefore more responsive. The pattern also appears in the authors' prior work studying the effects of

2. In future work, it would be helpful to confirm that the results regarding migration responses over time are robust to changing the level of geographic aggregation, acknowledging that the county-level data are only available starting in 1976.

unemployment changes on local net migration (Foschi and others 2023), suggesting that it is more than a statistical artifact of the present analysis.

What might explain this surprising pattern? In Foschi and others (2023), the authors speculate that shorter-distance moves are driven less by systematic economic conditions and more by idiosyncratic factors. One way to assess this argument in future work would be to restrict attention to moves motivated by economic considerations in surveys that record the primary reason for moving, such as the Current Population Survey (CPS) Annual Social and Economic Supplement (ASEC).[3] Another possibility is that correlated shocks to outside option locations may lead to stronger bias toward zero, as argued by Borusyak, Dix-Carneiro, and Kovak (2023), when employing finer geographic definitions. If shocks facing the focal location and relevant outside option locations are positively correlated, then estimated migration responses will be biased toward zero in comparison to a situation in which only the focal location faced the observed shock. In section II.A, the authors confirm that such bias is present in their analysis with state-level geography, though the magnitude of bias is small. In a footnote, they report that the bias is similar when using more detailed commuting zone or county-level geography, so this mechanism is unlikely to explain the differences across levels of aggregation.

Another potential explanation is that estimated effects are biased toward zero when the geography used in the empirical analysis is finer than the level at which economic equilibrium operates. To make this point clear, assume that states are the relevant local labor markets, such that the correct shock driving outcomes in the data generating process is a state-level Bartik shift share, $z_i^s \equiv \sum_i \frac{Emp_{si}^{t_0}}{Emp_s^{t_0}} d\ln Emp_i$, where s indexes states and i indexes industries. Also assume that the empirical analysis uses county-level geography, with a county-level Bartik shift share, $z_i^c \equiv \sum_i \frac{Emp_{ci}^{t_0}}{Emp_c^{t_0}} d\ln Emp_i$, where c indexes counties. By definition, the employment-weighted average of county-level Bartik shocks equals the state-level shock, that is, $z_i^s = \sum_{c \in s} \frac{Emp_c^{t_0}}{Emp_s^{t_0}} z_i^c$. Therefore, even if the county-level shocks are as good as random, the fact that the true shock driving outcomes is the state-level Bartik implies that a county-level analysis will be subject to attenuation bias.

3. Keep in mind that the limited sample size of the CPS makes substate migration analysis difficult, and note the important measurement issues raised by Kaplan and Schulhofer-Wohl (2012).

Note that this problem is isomorphic to classical measurement error but does not result from finite-sample noise in survey data; it would apply even if the county-level shocks were measured perfectly, without sampling error. The attenuation bias arises solely due to the use of a shock measured at an overly disaggregate geographic level. While in this example I assumed that the correct level of geographic aggregation was the state, in most cases states seem far too large to plausibly represent local labor markets. This argument is therefore most applicable to the comparison of commuting zone and county-level analyses, since commuting zones are plausibly good approximations of the relevant geographic scope of labor market equilibrium.

CONCLUSIONS This paper accomplishes an impressive feat of compiling seventy years of consistent data to facilitate the study of migration responses to labor demand shocks over a huge span of time. In doing so, it places scattered data points from prior literature in a consistent historical series and uses a variety of shock measures to bolster the credibility of the conclusion that migration remains an important mechanism underlying employment responses to local shocks. I find this conclusion entirely convincing and do not see it as inconsistent with the recent literature. The more provocative finding, that there has been no meaningful change in migration responsiveness over the past seventy years, is supported by the Bartik shift-share analysis alone. This conclusion also focuses on the migration elasticity, which I have argued is less relevant to arguments on declining dynamism than the direct effects of economic shocks on migration, since changes in migration costs affect both migration and employment responses to shocks. Despite these modest concerns and the suggestions above on adjustments to future empirical analyses, this paper is extremely valuable in introducing the literature to the idea of studying very long-term variation in migration responses. I expect that it will spawn many follow-up papers, which will in turn deepen our understanding of the role of migration in the dynamism of US labor markets.

REFERENCES FOR THE KOVAK COMMENT

Amior, Michael, and Alan Manning. 2018. "The Persistence of Local Joblessness." *American Economic Review* 108, no. 7: 1942–70.

Autor, David, David Dorn, and Gordon Hanson. 2021. "On the Persistence of the China Shock." *Brookings Papers on Economic Activity*, Fall: 381–447.

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

Blanchard, Olivier J., and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, no. 1: 1–75.

Borjas, George J. 2001. "Does Immigration Grease the Wheels of the Labor Market?" *Brookings Papers on Economic Activity*, no. 1: 69–133.

Borusyak, Kirill, Rafael Dix-Carneiro, and Brian K. Kovak. 2023. "Understanding Migration Responses to Local Shocks." Working Paper, June. https://www.andrew.cmu.edu/user/bkovak/borusyak_dixcarneiro_kovak.pdf.

Cadena, Brian C., and Brian K. Kovak. 2016. "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession." *American Economic Journal: Applied Economics* 8, no. 1: 257–90.

Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. "Regional Labor Market Adjustment in the United States: Trend and Cycle." *Review of Economics and Statistics* 99, no. 2: 243–57.

Davis, Steven J., R. Jason Faberman, and John Haltiwanger. 2012. "Labor Market Flows in the Cross Section and Over Time." *Journal of Monetary Economics* 59, no. 1: 1–18.

Ellison, Glenn, and Edward L. Glaeser. 1999. "The Geographic Concentration of Industry: Does Natural Advantage Explain Agglomeration?" *American Economic Review* 89, no. 2: 311–16.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar. 2023. "Labor Mobility and Unemployment over the Business Cycle." *AEA Papers and Proceedings* 113: 590–96.

Hyatt, Henry R., and James R. Spletzer. 2013. "The Recent Decline in Employment Dynamics." *IZA Journal of Labor Economics* 2: article no. 5.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2012. "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey." *Demography* 49, no. 3: 1061–74.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the US Labor Market." *Brookings Papers on Economic Activity*, Spring: 183–259.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25, no. 3: 173–96.

Rambachan, Ashesh, and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies* 90, no. 5: 2555–91.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127, no. 5: 2505–58.

COMMENT BY

ABIGAIL WOZNIAK[1] This paper examines the relationship between labor migration and economic shocks within the United States, offering a fresh perspective on a topic that has been extensively studied. Foschi, House,

1. Affiliation and disclaimer: The author is a vice president and senior economist at the Federal Reserve Bank of Minneapolis. Any opinions or conclusions expressed herein are those of the author and do not indicate concurrence with other members of the research staff of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the Board of Governors.

Proebsting, and Tesar contribute to this extensive literature in two ways. They go beyond the typical approach in past studies by expanding their empirical analysis along several dimensions. They also study a novel outcome of interest—a semi-elasticity of population change with respect to a local change in employment. The key finding of their research is observed stability in population responses to local employment shocks since the mid-twentieth century. At first glance, this stability contrasts with trends indicated by other data sources, which have often pointed toward a decline in migration rates. The authors argue that despite the general decline in migration observed through other sources, the responsiveness of migration to local economic shocks has remained constant.

A PAPER WITH MANY STRENGTHS AND AN IMPORTANT CONTRIBUTION This excellent paper has many strengths. First, it introduces a novel perspective to a mature area of study, raising useful questions about how to interpret the stylized facts around migration trends in the United States. In particular, the authors challenge readers to reconsider emerging evidence around the dynamics of migration, which finds that migration has become less responsive to local shocks over time (Partridge and others 2012; Dao, Furceri, and Loungani 2017). They ask whether this consensus aligns with results when the data on population changes are expressed as an elasticity as opposed to levels alone and when considered over a longer time period. This is particularly important in the context of a phenomenon like migration, which is influenced by a multitude of factors that may evolve over time.

Second, the paper's central finding is both intriguing and significant. The authors find that they cannot reject an elasticity of population with respect to jobs that is constant at about 0.5 over the six decades between the mid-1950s and the late 2010s. The authors' assertion that population responses to local employment shocks have remained largely stable over the past several decades at first appears in tension with the finding established in other work that migration has been on the decline over this period (Molloy, Smith, and Wozniak 2011). This tension warrants further exploration and has the potential to reshape our understanding of labor market adjustments. It suggests that while overall migration rates may have decreased, the propensity of individuals to move in response to economic shocks has not diminished.

A final strength of the paper is its comprehensive use of empirical data throughout. The authors draw upon a multitude of data sets and examine a wide array of dimensions, including various outcomes, geographical areas, time periods, and sources of local economic shocks. They have put all available sources on populations, migration flows, and local shocks to

thorough use. It can be difficult to present analysis from multiple sources across so many dimensions in a clear and coherent manner. The authors should be commended on accomplishing this.

This impressive paper builds on an extensive body of literature that is punctuated by a few key findings. I present an abridged overview here. The question of how local economies adjust to shocks was the focus of an influential *BPEA* paper by Blanchard and Katz in 1992. Their paper is known for its answer to this question. Blanchard and Katz (1992, p. 3) concluded that migration was the primary means of this adjustment, stating "a state typically returns to normal after an adverse shock not because employment picks up, but because workers leave the state." Lacking direct data on migration flows, they analyzed how populations change alongside local participation and unemployment rates in response to shocks affecting state-level labor markets. Their analysis provided a framework for understanding how labor markets respond to economic disruptions and how migration acts as a mechanism for restoring equilibrium. They emphasized that when a local economy experiences a downturn, migration can help to alleviate the negative impacts by allowing workers to move to areas with better opportunities.

Blanchard and Katz made use of contemporaneous work by Tim Bartik, whose 1991 book emphasized the importance of distinguishing between changes in labor supply and changes in labor demand when examining local outcomes. Bartik underscored the need for researchers to carefully identify underlying causes of changes in local employment in order to accurately assess the connections between employment and other local outcomes. He argued that it is essential to differentiate between situations where employment declines due to a decrease in demand for labor and those where it declines due to a decrease in the supply of labor. His approach to resolving this identification challenge—a shift-share measure of local employment growth changes—was adopted by Blanchard and Katz and has since become a cornerstone of the literature on local economic adjustment.

The conclusions in Blanchard and Katz (1992) were revisited in a 2017 paper by Dao, Furceri, and Loungani. By revisiting the analysis conducted by Blanchard and Katz with more detailed and up-to-date data, including direct measures of migration, Dao, Furceri, and Loungani's (2017) research suggested that migration's role in re-equilibrating local labor markets was less significant than initially estimated by Blanchard and Katz. Their work also indicated that the responsiveness of migration to local shocks had declined substantially in recent years, challenging the earlier understanding of migration as a key adjustment mechanism. This finding raised concerns

about the ability of labor markets to adjust to economic shocks through migration.

Between Blanchard and Katz (1992) and the update by Dao, Furceri, and Loungani (2017), papers appeared first documenting and then investigating a multidecade decline in migration. Molloy, Smith, and Wozniak (2011) identified a steady decline in migration in the United States beginning around 1980. This finding, observed across a variety of data sources, spurred further research into the causes and consequences of this decline. While several potential explanations have been proposed, including factors such as increased housing costs, greater attachment to place, and changes in the labor market, a significant portion of the decline remains unexplained, suggesting that the factors influencing migration patterns are complex and multifaceted. Some studies have suggested a link between this decline in migration and changes in the labor market, such as the decline in job-to-job transitions documented by Molloy, Smith, and Wozniak (2017), but the evidence supporting this connection is not definitively causal.[2] The relationship between labor market dynamics and migration patterns remains an area of active investigation.

Foschi, House, Proebsting, and Tesar make a significant contribution to this ongoing discussion by presenting evidence that migration has not declined when measured as an elasticity with respect to employment changes. The authors' central argument revolves around an analysis of the relationship between changes in log jobs and log population in response to local shocks to employment growth, as illustrated in the paper's figure 5. To arrive at this elasticity, the authors estimate the percentage change in local jobs five years following a Bartik-style shock (figure 5, panel A) and the percentage change in population following the same shock (panel B). They take the ratio as their preferred measure of the elasticity of population with respect to jobs (panel C) and generate confidence intervals around the long time series of elasticity estimates using the delta method to estimate standard errors. This analysis reveals that the ratio of population change to employment change has exhibited remarkable stability over time, with an estimated elasticity of approximately 0.5. This finding suggests that, historically, population adjustments in response to employment shocks have been consistent, challenging the notion that migration has become less responsive to economic conditions. The authors interpret this stability

2. There is also evidence for other explanations. Kaplan and Schulhofer-Wohl (2017) present evidence in favor of a convergence in regional labor market opportunities and reduced experimentation with local amenities. Olney and Thompson (2024) present new evidence for the importance of the housing market channel.

as evidence that labor markets continue to adjust to shocks in a predictable manner.

They propose that the moderated findings in some previous studies, such as those in Dao, Furceri, and Loungani (2017) may be attributable to a period characterized by small shocks as measured by employment changes. They argue that when employment changes are small, the corresponding population changes will also be small, which could lead to the appearance of declining responsiveness. Their paper offers a new lens through which to interpret migration patterns and their relationship with economic shocks, suggesting that the underlying responsiveness of migration to economic conditions may be quite stable.

HOW ROBUST IS THE PAPER'S CENTRAL FINDING? The authors put in excellent effort to construct a reliable long-run measure of this population elasticity, and their interpretation of a stable elasticity of 0.5 is one reasonable way to read the fitted line in figure 5, panel C. However, there is a reasonable alternative interpretation, discussed in the next section, and it is worth asking more generally how robust this central finding is.

The first cause for concern is the presence of large standard errors associated with the elasticity estimates. These imply a substantial degree of uncertainty in the estimates, such that the annual estimates are not statistically different from zero about one-third of the time. This lack of precision raises questions about the reliability of the 0.5 elasticity estimate as a fixed parameter. It leaves open the possibility that the true relationship between employment shocks and population change may vary more than the point estimate suggests.

The presence of two types of measurement issues also raises concerns about whether estimates are comparable across the entire time series. The first issue is that a main object of interest—population—is not measured directly in most years in the data. Instead, it is estimated using an interpolation between decadal observations. The use of interpolated decennial population data may artificially smooth underlying annual population changes. Interpolation techniques can mask short-term fluctuations in population, which may be relevant to understanding the dynamic relationship between employment shocks and migration. Investigating the implications of this smoothing for the relationships in the time series is beyond the scope of this discussion, but intuitively it would seem to have the potential to introduce a mechanical relationship that contributes either to stability in the estimates, downward-biased standard errors (via noise reduction), or both.

A second issue is that most series used by the authors—in figure 5 as well as other estimates—contain seams at which either closely related

series are spliced together, such as the Internal Revenue Service (IRS) data, or classifications that change meaningfully across adjacent years, such as changes in geographic aggregation and industry codes. These seams raise concerns about the comparability of the data over the extended six-decade period examined in the study. Ensuring data consistency over such a long time frame is a significant challenge, and these seams may affect the validity of the conclusions drawn.

In addition to the concerns about data quality and estimation precision, the authors' empirical results present some puzzles. The first puzzle is the finding that migration responses appear to be weaker at finer geographic levels, as shown in the paper's table 1. This result is counterintuitive, as one might expect individuals to be more responsive to local shocks when the distance required to move is smaller. The authors acknowledge this unexpected pattern but do not provide an explanation. It is possible that this discrepancy arises from issues related to data construction, particularly the allocation of employment data at finer levels of geography. The challenges of accurately measuring employment and migration flows at the county level, for example, may contribute to this unexpected finding. However, these in turn raise questions about which geographic scale is most accurate and also whether the factors that influence migration decisions differ across different geographies.

A second puzzle relates to the authors' findings on migration responsiveness across different age groups reported in the paper's figure 6. As expected, the authors find that migration responsiveness declines with age. However, the results also show that responsiveness is surprisingly similar for workers between the ages of 30 and 64, with a significant drop-off only occurring after age 65. The similar responsiveness of workers across these ages is unexpected, as one might anticipate a more gradual decline in responsiveness with age.

Resolving these puzzles more satisfactorily would allow readers to have greater confidence in the overall conclusion of a single, stable, long-run population elasticity.

DIRECTIONS FOR FUTURE RESEARCH This section offers suggestions for future research. I discuss ways to improve the evidence in the paper and then turn to bigger questions that might be the focus of new research projects.

To address the concerns about data quality and to further demonstrate the robustness of their central finding, the authors could identify the highest-quality data environment available in their study and emphasize this as a related preferred estimate. Rather than relying on the longest possible series at all times, this analysis would rely on the highest-quality data scenario for

measuring local employment changes and local population changes repeatedly over time. This may require focusing on specific time periods and data sources, such as the IRS data series on migration flows prior to 2012, which are considered to be more reliable. The potential benefit is that finding a similar population elasticity in this higher-quality data context would bolster the authors' argument for an invariant elasticity. Additionally, researchers should consider incorporating other types of shocks into their analyses, including those related to export employment, the North American Free Trade Agreement (NAFTA), and the Great Recession. These alternative shocks may provide valuable insights into the heterogeneity of migration responses to different types of economic disruptions as well as responding to concerns raised in the companion discussion by Brian Kovak.

I also suggest the authors also take more seriously an alternative interpretation of the findings presented in the paper, namely, that the relationship between population adjustment and employment change is not constant over time but instead varies across different periods. The data presented in figure 5, panel C, also appear consistent with this alternative view. They suggest a cyclical pattern, with periods of stronger and weaker adjustment. These fluctuations may stem from the impact of the initial shocks, as figure 5, panel A, shows periods in which jobs change relatively little following a shock, or from changes in responsiveness more generally. The authors attribute smaller population changes following shocks in these periods to the fact of smaller employment changes. But this begs the question of why shocks generate bigger or smaller employment changes in different periods. It deserves to be considered as an important question in this analysis. Future research could investigate factors that drive this variation, such as changes in economic conditions, policy interventions, or structural shifts in the labor market. Understanding these dynamic relationships is crucial for developing a comprehensive model of migration behavior and could help to explain why migration patterns appear to have changed over time.

The paper raises several important questions that warrant investigation in future research. As already noted, figure 5, panel A, suggests that the employment response to shocks exhibits substantial fluctuations. Future research may want to take up the question of what drives this. A deeper understanding of the mechanisms that influence the magnitude of employment changes following an economic disruption is essential for accurately interpreting the relationship between employment changes and migration patterns, as in this paper, but it potentially has implications for our understanding of local adjustment more broadly. Researchers should explore the role of industry structure, labor market institutions, and regional

Figure 1. Percent Moving, Age 16 and Older

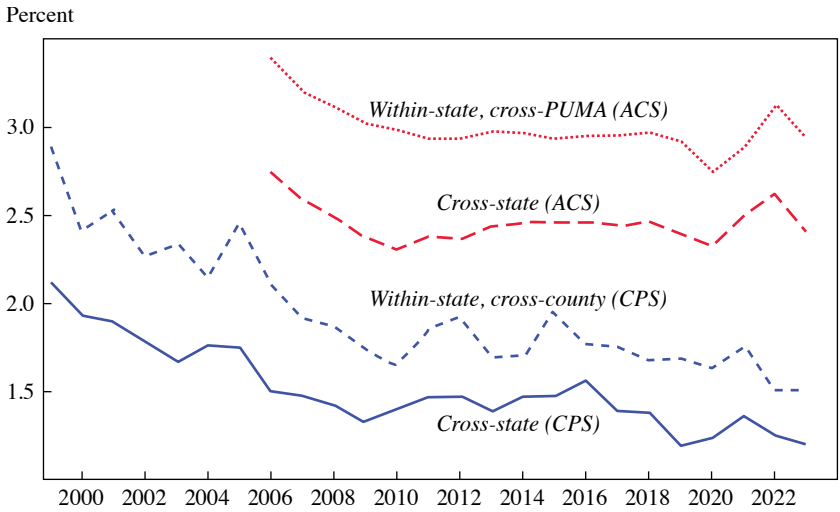

Source: Christopher L. Smith's calculations from the Current Population Survey (CPS) and American Community Survey (ACS), following and extending the approach in Molloy, Smith, and Wozniak (2011). PUMA = Public Use Microdata Area.

characteristics in shaping these responses. Additionally, it would be valuable to examine how factors such as population characteristics, unemployment benefits, job training programs, and housing policies affect the speed and magnitude of employment adjustments.

Another important question concerns the implications of stable migration responsiveness in the context of declining overall migration rates. Many papers in the literature study the relative roles of two broad determinants to the migration decline: dispersion in economic opportunity and in amenities (Leduc and Wilson 2024; Kaplan and Schulhofer-Wohl 2017; Partridge and others 2012). If individuals and households are just as responsive to local economic shocks as they have been in the past, but overall migration rates have declined (as shown in figure 1 here), then migration driven by shocks may now constitute a larger share of total migration.

If the composition of migrants has shifted toward those responding to shocks, this may be further evidence of amenity equalization (via prices) across the United States, as suggested in Leduc and Wilson (2024). On the other hand, Molloy and others (2016) and Molloy, Smith, and Wozniak (2017) provide evidence linking the decline in migration to labor market fluidity. If workers continue to relocate in response to local shocks but face

a less fluid labor market when they do so, this reallocation may be costly to workers and less efficient for firms. Equalized amenities across space could provide some insurance against a welfare-reducing move, but the overall welfare implications of this constellation of changes are difficult to intuit. Future research should pursue a framework that can study net welfare changes in a realistic setting that accommodates these various trends.

In conclusion, this paper makes a valuable contribution to the literature on labor migration and economic shocks in the United States. The finding of stable population responsiveness to local employment shocks is an important complement to previous research and opens up new avenues for investigation. However, the puzzles in the empirical results, questions surrounding data quality, and overall large standard errors suggest caution in applying the paper's main estimate as a fixed parameter across different periods and contexts. Future research should address these concerns by employing high-quality data, exploring alternative interpretations, and focusing on the key questions raised by this paper. A more nuanced understanding of the factors that influence migration decisions, including the role of economic shocks, is crucial for informing policies aimed at promoting labor market efficiency and regional economic development. It is essential to consider the broader context of declining migration rates and to investigate the reasons behind this decline in order to fully understand the implications of the authors' findings.

REFERENCES FOR THE WOZNIAK COMMENT

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

Blanchard, Olivier J., and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, no. 1: 1–75.

Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. "Regional Labor Market Adjustment in the United States: Trend and Cycle." *Review of Economics and Statistics* 99, no. 2: 243–57.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration." *International Economic Review* 58, no. 1: 57–94.

Leduc, Sylvain, and Daniel Wilson. 2024. "Snow Belt to Sun Belt Migration: End of an Era?" Working Paper 2024–21. San Francisco: Federal Reserve Bank of San Francisco. https://doi.org/10.24148/wp2024-21.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers on Economic Activity*, Spring: 183–237.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25, no. 3: 173–96.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2017. "Job Changing and the Decline in Long-Distance Migration in the United States." *Demography* 54, no. 2: 631–53.

Olney, William W., and Owen Thompson. 2024. "The Determinants of Declining Internal Migration." Working Paper 32123. Cambridge, Mass.: National Bureau of Economic Research.

Partridge, Mark D., Dan S. Rickman, M. Rose Olfert, and Kamar Ali. 2012. "Dwindling U.S. Internal Migration: Evidence of Spatial Equilibrium or Structural Shifts in Local Labor Markets?" *Regional Science and Urban Economics* 42, no. 1–2: 375–88.

GENERAL DISCUSSION Many attendees discussed potential interpretations of the authors' results, particularly that labor migration elasticity, denoted by gamma in the research, is constant over the period of interest. Steven Davis noted that the standard errors around gamma are very large. He argued that even though one cannot reject the hypothesis that gamma has remained roughly constant over time, they also cannot reject other hypotheses that fit within the standard errors. Davis proposed, for example, that gamma likely depends on the age distribution of the affected labor supply, since older workers are less likely to relocate in response to idiosyncratic shocks. He asked the authors to explain how their findings are consistent with the large shifts in the age distribution during the period of interest.

In response, Christian Proebsting described an exercise the authors conducted to investigate this issue. He agreed that their estimates of gamma vary with respect to age. However, their exercise indicated that applying these coefficients to composition changes over time does not result in significant enough differences to cause measurable trends in the aggregate.

Davis also proposed that gamma may vary with respect to the size of the local demand shock, arguing that workers who are close to indifferent between moving and staying would be enticed to move given a larger shock. This argument further suggests that gamma is connected to cyclical variation in the macroeconomy and heterogeneous across regions. He conjectured that both alternatives indicate the data could be better explained by a benchmark population-level gamma, which may vary with respect to time and exhibit localized variations due to regional demographic trends, labor market shocks, and heterogeneous macroeconomic effects.

Stan Veuger proposed that gamma may respond to the cost of moving, including housing prices, which could also result in a cyclical trend. He observed that the authors' model considers housing prices as a component

of household wealth but not cost of moving. He suggested that the authors include the latter factor in their model as recent studies have demonstrated the importance of housing prices for labor mobility.

Tara Sinclair concurred in cautioning against the conclusion that gamma has remained constant. She pointed out that such a strong conclusion could have significant implications for future researchers if they assume in their models that labor reallocation rates are constant over time.

Linda Tesar acknowledged that this conclusion needs greater scrutiny, and that gamma may have a cyclical component. She explained that their motivation for the paper was to dispel the increasingly common assumption in trade geography literature that labor does not move between markets but capital does. Therefore, it was encouraging to her that their findings show gamma has not been dramatically falling, even if they do not have sufficient evidence that it is constant.

John Haltiwanger contrasted the authors' results with findings in his 2014 paper coauthored with Davis.[1] They found that the worker reallocation rate, defined similarly to gamma, has been declining over time. Their paper suggests that the effect is driven by decreasing pace of reallocation among younger and less educated workers, coupled with an aging workforce where older workers have consistently lower pace of reallocation. They find this decline pervasive within geographic areas. They emphasized that there are adverse effects of this decline in terms of the ability of young and less educated workers to remain attached to the labor force and to move up the job ladder. Related research finds decline in the pace of dynamism (business expansion, contraction, entry and exit), which is associated with a decline in the responsiveness of firms to idiosyncratic productivity shocks.[2] The point of these two examples is to highlight that it is important to consider the potential consequences on key outcomes of changes in the pace of the dynamics of the labor market.

Joe Beaulieu also expressed his surprise that gamma is constant and proposed a potential econometric explanation for this result. He suggested that in periods with few national labor market shocks, there is very little variation in gamma. If one adds some random error over time, then these periods with low variation will have very high standard errors. Proebsting

1. Steven J. Davis and John Haltiwanger, "Labor Market Fluidity and Economic Performance," working paper 20479 (Cambridge, Mass.: National Bureau of Economic Research, 2014).

2. Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Changing Business Dynamism and Productivity: Shocks versus Responsiveness," *American Economic Review* 110, no. 12 (2020): 3952–90.

provided a similar econometric explanation: Because gamma is a ratio whose denominator is local change in employment, standard errors increase rapidly as the change in employment approaches zero.

Jonathan Pingle remarked that apart from gamma, the authors captured a decline in gross labor migration that has been ongoing for seventy-five years. He elaborated that prior research has attempted to explain this trend with changes in labor supply, including demographics, retirement, education, marriage, and children, but supply is difficult to identify separately from changes in labor demand. He applauded the authors for successfully disentangling labor supply and demand by employing the Internal Revenue Service (IRS) data, thus achieving what he believed to be the greatest contribution of this paper to the literature.

Jonathan Wallen asked the authors to reconcile this decline in gross migration with their main result that gamma has remained constant. He conjectured that firms may be less geographically sensitive to labor demand over time. Christopher House provided an alternative explanation utilizing a Bartik instrument, observing that the responsiveness to Bartik shocks has been growing over time, but the frequency of these shocks has decreased.

Anil Kashyap inquired about the mechanisms that led to geographic and socioeconomic heterogeneity in the authors' results. To authors' finding that rural and low-income areas have lower mobility, he wondered whether those areas were disproportionately impacted by the China shock due to the offshoring of manufacturing and low-skill jobs. He then asked what type of shock would have the greatest welfare impact, and which areas or populations would be most vulnerable to such a shock due to income, skill bias, and technical change. Tristan Reed questioned the authors' result that rural areas have lower mobility than urban areas, given evidence in the literature that the latter has inelastic housing supply and housing costs are a significant barrier to moving. He suggested estimating the gamma coefficient separately for each state and comparing those with greater urban populations, for example, New York, to those with mostly rural populations, for example, Texas.

Several attendees commented on the authors' empirical methods. Katharine Abraham remarked that, due to their small size, counties are less suitable as the geographic unit for modeling labor market migration than commuting zones, which are defined specifically based on commuting patterns. In response, Andrea Foschi referred to their finding that only a quarter of labor migration occurs between counties within the same state after conditioning on labor demand shocks, meaning the vast majority of mobility due to labor market and economic conditions occurs between states.

David Romer encouraged the authors to derive their empirical specification by first using a simple theoretical model of an outcome variable, such as population level. He mentioned the paper by Jacob Orchard, Valerie Ramey, and Johannes Wieland who conducted a similar exercise and found that most leading estimates of the marginal propensity to consume from the 2008 financial crisis are biased due to model misspecification.[3] He suggested that the authors' shocks are likely serially correlated, so this exercise might provide insights that local projection regressions could not give.

Some commenters also expressed concerns regarding the paper's policy implications. Abraham noted that the authors' findings might seem to imply that, since workers often just relocate in response to labor market shocks, place-based policies may be unnecessary to support employment in struggling geographies. From a welfare perspective, however, the types of jobs, as well as wages and hours in the adversely affected geographies also are important. While some people may be able to emigrate following a negative shock to another area, for a variety of reasons, many will stay put. If the jobs available to the stayers are of lower quality, that represents a decline in welfare.

Haltiwanger also proposed some implications for welfare and productivity from labor migration. He referred to the aforementioned paper he coauthored with Davis, which finds declining dynamism in the retail sector may be caused by mom-and-pop businesses being replaced by large national chains, which have notable productivity gains. However, lower mobility for younger and less educated workers implies that they may have greater difficulty advancing in their career. He suggested the authors conduct a similar analysis to understand whether their results contain implications for wages, hours, productivity, promotions, and welfare.

House clarified that their policy suggestions were not intended to imply that there are not significant effects of local labor shocks, such as the China shock, but simply to point out that some portion of the displaced workforce is capable and often does move to areas with better opportunities. He noted that this result is more encouraging than the alternative, where labor mobility has declined and therefore people are less capable of migrating to

3. Jacob D. Orchard, Valerie A. Ramey, Johannes F. Wieland, "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," *Quarterly Journal of Economics* 140, no. 3 (2025): 2001–52.

expanding labor markets or leaving contracting markets. Foschi added that their results show in-migration tends to respond more than out-migration, suggesting workers are more likely to move to expanding labor markets. This finding may also indicate that some workers can remain stuck after experiencing a negative labor shock, and place-based policies and welfare analyses may therefore remain important.

Online Appendix

Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks.

Andrea Foschi

Christopher L. House

Christian Proebsting

Linda L. Tesar

I Additional figures

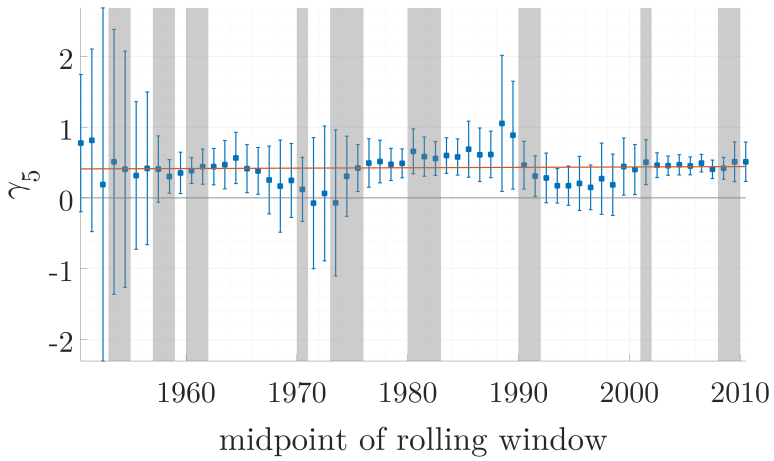

Figure 1: Ratio of population and jobs responses (γ_5^{POP})

Note: The figure plots the γ_5 coefficient defined obtained by taking the ratio of the β_5 coefficients for population and jobs, overlapped with NBER recessions (highlighted in grey). See notes to Fig. 5 for more details. *Source:* authors' calculations.

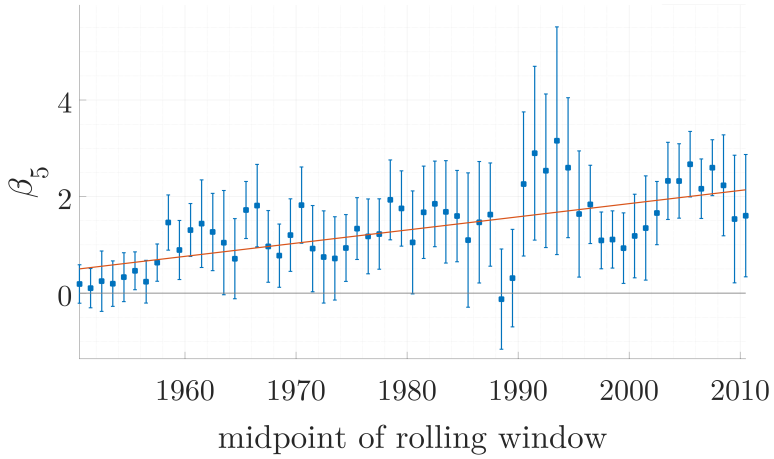

Figure 2: Response of labor market variables other than population ($\beta_5^{Jobs-Pop}$)

Note: The figure plots the estimate of the β_5 coefficient for the difference between jobs growth and population growth (which captures the growth in jobs due do other labor market variables, i.e. jobs/employment ratio, participation and unemployment). See notes to Fig. 5 for more details. *Source:* authors' calculations.

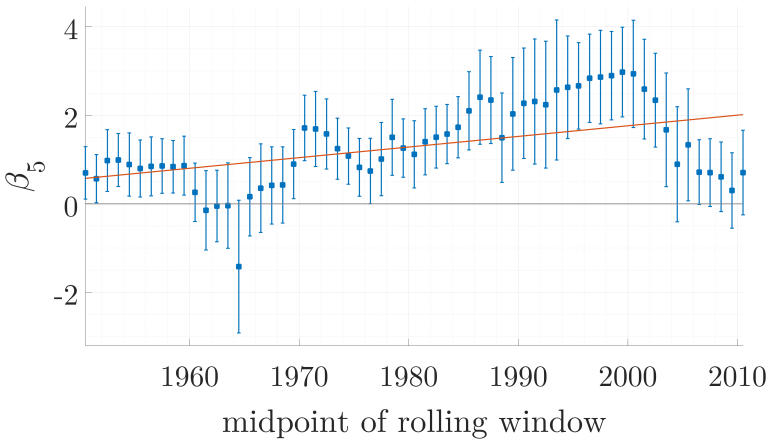

Figure 3: WAGE RESPONSE (β_5^{Wage})

Note: The figure plots the estimate of the β_5 coefficient for wages ("wages and salaries" from the BEA divided by the number of jobs). See notes to Fig. 5 for more details. *Source:* authors' calculations.

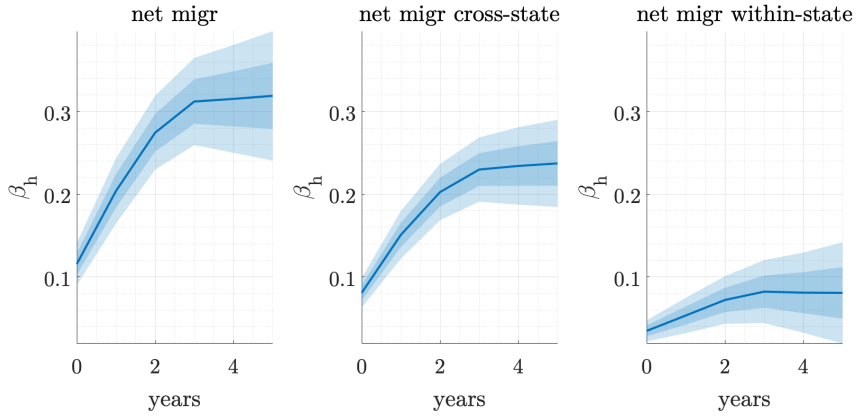

Figure 4: CROSS-STATE VS WITHIN-STATE MIGRATION RESPONSES

Note: The figure plots the estimated β_h coefficients obtained for each migration measure at different horizons h (x-axis). The sample is at the county level and runs from 1976 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% and 95% confidence intervals. *Source:* authors' calculations.

II Additional tables

Corr($\gamma_{5,t}, \%\Delta GDP_{t+j}$)										
$j=-5$	$j=-4$	$j=-3$	$j=-2$	$j=-1$	$j=0$	$j=1$	$j=2$	$j=3$	$j=4$	$j=5$
-0.102	0.021	0.005	-0.012	0.058	0.095	0.043	0.037	-0.109	-0.029	-0.100

Table 1: Lead-lag correlations between γ_5 and GDP

Note: The table reports the lead-lag correlations between the γ_5 (as reported in Fig. 5) and national real GDP growth (from the FRED). *Source:* authors' calculations.

		Short-run ($h=0$)			Long-run ($h=5$)		
		ur	Lfp	Pop	ur	Lfp	Pop
γ	State	-0.80	0.03	0.16	-0.21	0.12	0.67
		(0.13)	(0.14)	(0.10)	(0.07)	(0.05)	(0.08)
	CZ	-0.45	0.44	0.11	0.07	0.44	0.63
		(0.04)	(0.04)	(0.03)	(0.08)	(0.07)	(0.08)
	County	-0.39	0.49	0.12	0.06	0.61	0.45
		(0.04)	(0.04)	(0.02)	(0.05)	(0.06)	(0.07)
$\gamma^{partner}$	State	-0.14	0.10	-0.24	-0.50	0.69	-1.19
		(0.73)	(0.87)	(0.53)	(0.44)	(0.51)	(0.67)
	CZ	-0.71	-0.70	-0.01	-0.64	-0.38	-0.26
		(0.14)	(0.13)	(0.07)	(0.16)	(0.16)	(0.20)
	County	-0.61	-0.54	-0.07	-0.44	-0.60	0.16
		(0.10)	(0.10)	(0.06)	(0.09)	(0.13)	(0.14)

Table 2: Outside options corrections at different levels of aggregation

Note: The top half of the table displays the estimates of the γ_h coefficients obtained by running the outside-options-corrected regression (7) 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 $\gamma_h^{partner}$ coefficients from the same regression. The sample period is from 1976 till 2016, with projections going up to 2021 (for $h=5$). Standard errors (in parentheses) are clustered at the level of the region. All regressions are weighted by population. *Source:* authors' calculations.