Inter-state Labor Mobility and the U.S. Economy *

Andrea Foschi
Bank of Italy

Christopher L. House
University of Michigan and NBER

Christian Proebsting
KU Leuven

Linda L. Tesar
University of Michigan and NBER

August 7, 2025

bstract

The aim of this paper is to summarize the facts on changing U.S. labor mobility and the factors driving its decline. We draw four conclusions from the data. First, the decline in gross inter-state migration over the last 50 years is relatively modest and has been essentially stable since the early 2000's. Second, the *net* migration rate, which is one of the main equilibrating mechanisms between locations, exhibits no trend. Our estimate of the elasticity of net migration to regional labor demand shocks is about 0.6 and has remained essentially constant since the 1950s. Third, there is little evidence that demographic changes explain the decline in aggregate gross migration. Fourth, we find that there has been a decline in the cross-state dispersion in the labor demand shocks that generate movements in population from one state to another. This decline reflects both the fact that the industry composition of states has become more similar, and that the shocks to industries have become more correlated. This finding could be related to other evidence that the factors that drive turnover or churn in labor markets has declined.

JEL Codes: E24, E32, F66, J61, R23

*We gratefully acknowledge financial support from the Michigan Institute for Teaching and Research in Economics (MITRE). We thank Emmanuel Boah and Erin Deal for their excellent research assistance. We thank Gaetano Basso for useful comments. Foschi: afoschi@umich.edu; House: chouse@umich.edu; Proebsting: Christian.Proebsting@kuleuven.be; Tesar: ltesar@umich.edu. The views expressed here are those of the authors and do not necessarily reflect those of the Bank of Italy or the Eurosystem. *Note: This paper is prepared for the August 2025 Jackson Hole Conference on Labor Markets in Transition: Demographics, Productivity, and Macroeconomic Policy.*

1 Introduction

In 1978, roughly 3 out of every 100 people in the U.S. moved across states each year. Today, that rate is closer to 2.5 percent – reflecting a growing reluctance on the part of Americans to move. The responsiveness of the labor force to economic opportunities across geographic locations is considered to be a key contributor to the fluidity and dynamism of the U.S. labor market. The decline in labor mobility has fueled concerns that this engine of U.S. growth and efficiency is slowing. The change in migration rates has occurred together with other trends affecting the U.S. labor market including the aging of the population, higher average levels of education, a greater fraction of foreign-born citizens as well as shifts in the industrial composition of the economy.

The aim of this paper is to evaluate what we know, or what we *think* we know, about the gradual change in labor mobility in the United States over the past 50 years and its implications for the macroeconomy. We draw four conclusions from the data. First, the decline in gross inter-state migration over the last 50 years has been relatively modest and the level of mobility has essentially stabilized since the early 2000's. This conclusion is based on data from the IRS, which, for reasons we discuss below, we think is the most reliable measure of aggregate labor mobility in the United States. Second, the *net* migration rate, which is one of the main equilibrating mechanisms between locations, exhibits no trend. Our estimate of the elasticity of net migration to regional labor demand shocks is about 0.6, meaning that for a one percent increase in labor demand, roughly 60 percent of the equilibrium change in employment will be met through a net inflow of workers from a different state. Furthermore, we find that the estimate of the net migration elasticity has remained essentially constant since the 1950s. Third, confirming what much of the literature on migration has found, there is little evidence that demographic changes explain the decline in aggregate gross migration. The U.S. population is indeed older, more educated and more likely to be foreign born. Those groups have different migration propensities, but when taken together, these shifts have no net explanatory power for total migration.[1] Fourth, we find that there has been a decline in the cross-state dispersion in the labor demand shocks that generate movements in population from one state to another. This decline reflects both the fact that the industry composition of states has become more similar, and that the shocks to industries have become more correlated. This finding could be related to other evidence that the factors that drive turnover or churn in labor markets has declined. We leave the analysis of the connections between labor market turnover, business innovation and migration flows for future study.

[1] See, for example, Molloy et al. (2011) and Azzopardi et al. (2020)

We begin our analysis by considering the change in the gross migration rate – a measure of the total number of individuals that relocate across state borders as a share of the population. Many of the differences in how researchers characterize the change in the migration rate can be traced to differences in the data used to measure labor mobility. Our preferred measure is based on IRS tax return data that captures changes in location for all U.S. tax payers and is used by the Census to track the U.S. population. This measure shows a clear but modest decline of roughly 1/2 of a percentage point (from 3 to 2.5 percent) between 1976 and 2022 with a slight increase since 2010. We also consider two other common measures for labor mobility. The American Community Survey (ACS) asks respondents directly whether they have moved in the past year. This data, available only from the early 2000's, shows systematically lower migration rates than the IRS data but is nevertheless highly correlated with the IRS measures. We also consider data from the Current Population Survey (CPS). The CPS data shows the most dramatic decline in migration rates with less sign of leveling off. We will draw on the CPS to examine demographic factors that may contribute to changing patterns in US migration.

The decline in gross migration is largely accounted for by a decline in offsetting migration – the relocation of individuals across state borders that leave the state's total population unchanged. Offsetting migration flows account for roughly 90 percent of gross migration. We find a decline in offsetting migration that mirrors the pattern in gross migration, leveling off in the early 2000s. One interpretation of the decline in offsetting migration is a decline in turnover or churn as workers shift from location to location. The drop in offsetting migration from the 1980s to the 2000s coincides with the decline in the job reallocation rate and other measures of labor market fluidity studied by Decker et al. (2016, 2017) and Davis and Haltiwanger (2014). Molloy et al. (2017) also observe a secular decline in job changing and link it to the decline in gross migration. Interestingly, the surge in entrepreneurial business creation that Decker and Haltiwanger (2023) observe during the COVID period coincides with the uptick in migration that we see at the end of our sample.

The second component of gross migration is net migration across state borders. Despite its small fraction of total cross-state migration, net migration continues to play an important role in mitigating differences in economic outcomes across locations. In Section 3 we examine the response of net migration to local shocks, which we identify using an industry-composition (Bartik 1993) instrument for shifts in labor demand (see Foschi et al. forthcoming). Similar to the early results of Blanchard and Katz (1992), we find that, on impact, most of the local response to the shock is a reduction in unemployment, with only a small fraction that can be attributed to net in-migration. After five years, this pattern changes, and most of the response to the shock is an inflow of migration. We use the Bartik instrument to estimate an elasticity of net migration; essentially the ratio of the migration response to the employment response. Focusing on the five-year horizon, we find a migration elasticity of roughly 0.6. This elasticity is fairly constant over long samples of the data, going back to the 1950s. Further, an elasticity of roughly 0.6 emerges based on different levels of aggregation (state, commuting zone and countries) and for different shocks to labor demand.

Section 4 examines how shocks to labor demand have changed over time. We find that the dispersion in the Bartik instrument across states has declined since the 1980s and the decline in the dispersion is consistent with the decline in gross and offsetting migration rates. A decomposition of the Bartik suggests that about half of the decline in the dispersion of the shocks can be attributed to changes in the shifts and half to changes in the shares. In other words, industries are growing at more similar rates and the industrial composition across states has become more similar. These findings are consistent with the view that business cycles in the U.S. may have become more synchronized (for a recent contribution to this literature see Fieldhouse et al. (2024)).

2 Trends in U.S. Migration Rates

Figure 1 shows the gross cross-state migration rate in the U.S. based on data from three different sources. The gross migration rate is calculated as follows:

$$\text{Gross Migration Rate}_t = \frac{1}{2}\frac{\sum_i [\text{in-migrants}_{i\,t} + \text{out-migrants}_{i\,t}]}{\sum_i \text{population}_{i\,t}}.$$

We divide by 2 because every in-migrant to a state is also an out-migrant from another state – which leads to "double counting" migrants. Gross migration reflects the average number of individuals that changed their location from one state to another in a given year as a share of the population. The dark black line in the figure is based on data from the IRS. The IRS bases this measure on the number of tax returns with a change in filing address from the previous year. The green dashed line is based on data from the American Community Survey, which asks survey respondents whether they moved their residence from one state to another over the past year. In recent years about 3 million households participated in the ACS. Data from the IRS and the ACS are used to inform Census estimates of the U.S. population and its geographic distribution. The dotted blue line shows the migration rate based on data from the Current Population Survey

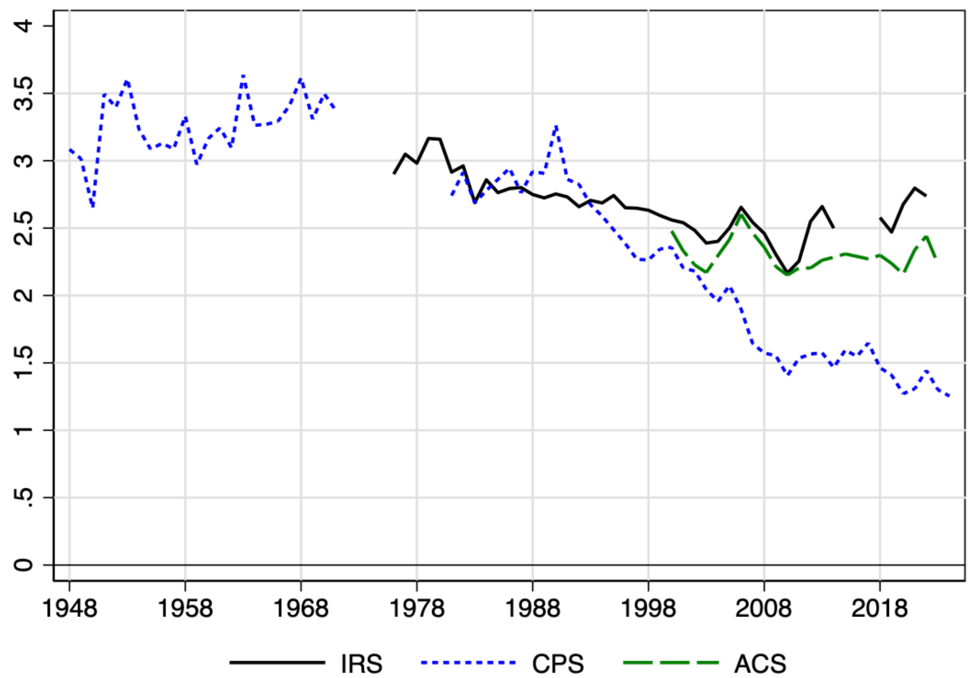

Figure 1: Gross migration rates

Note: The figure plots gross migration rates from three different sources: the CPS, the IRS and the ACS. Figures are expressed in percent of the US population and only refer to internal migration. There was a change in the method for collecting and producing IRS migration data, which generated large variations in the IRS figures for 2014-2016 (see DeWaard et al. 2022 for a complete discussion). We dropped those observations from the plot and used the three-year average instead. Including values for those years does not change the estimated trends.

(CPS) with a coverage of approximately 100,000 households.

The IRS data suggest that the migration rate declined from a peak of just above 3 percent in the late 1970s to a minimum rate of just above 2 percent in 2010. Thereafter the migration rate increased reaching 2.7 percent by the end of our sample. The migration rate based on the ACS sample is slightly below the IRS rate and is essentially flat since 2000. The CPS data display the most dramatic decline in gross migration. According to the CPS, the gross migration rate was above 3 percent in the 1950s and has steadily declined since then. The decline in the migration rate slowed after 2008, with a rate of 1.2 percent at the end of the sample.

U.S. residents are not required to register their local address and therefore it is not surprising that it is difficult to accurately track population movements over time. Each of the data sources that we and other researchers commonly use to study labor migration has advantages and disadvantages. The IRS data has the largest sample size as it includes all tax filers each year. A potential bias in the IRS data, however, is that the migration behavior of tax filers could differ from individuals who do not file taxes. Because the focus of our study is on the link between location decisions and labor market outcomes, the bias toward individuals who earn taxable income may be less of a concern. The IRS data do not contain information about the individual tax filer, so we must rely on the ACS and the CPS to study the characteristics of those who migrate. The fact that the migration rates based on the IRS and the ACS are highly correlated is reassuring and we place a heavier rate on these gross rates than those of the CPS.[2] The CPS, however, has the advantage of reporting detailed information about individual survey respondents over a long period of time, and we will draw on some of this information in the sections that follow.

Many economists and demographers have noted the widening gap between the migration rates based on the CPS and the rate based on the ACS/IRS data. Among the concerns about the CPS is the increased difficulty of reaching survey respondents over time, the fact that the CPS excludes some types of individuals that are picked up in the ACS, and misreporting. Matching individuals across datasets, economists from the U.S. Census show that a substantial fraction of documented cross-state movers in the administrative records data are not reported as movers in the CPS (Hyatt et al. 2018). This fraction has increased over time, and the divergence has become sufficiently concerning that the analysts at the Census suggest caution when using CPS data to study migration trends and recommend using IRS data instead.[3]

2.1 Decomposition of Gross Inter-State Migration

Gross migration can be decomposed into two types of labor flow. Since for any two numbers a and b we have $\min\{a, b\} = \frac{1}{2}[a + b \quad |a \quad b|]$ we can write the Gross Migration Rate as the sum of the Absolute Net Migration Rate and the Offsetting Migration Rate where

$$\text{Absolute Net Migration Rate}_t = \frac{1}{2}\frac{\sum_i |\text{in-migrants}_{i\,t} \quad \text{out-migrants}_{i\,t}|}{\sum_i \text{population}_{i\,t}}$$

and,

$$\text{Offsetting Migration Rate}_t = \frac{\sum_i \min(\text{in-migrants}_{i\,t}, \text{out-migrants}_{i\,t})}{\sum_i \text{population}_{i\,t}}.$$

Absolute Net Migration measures the extent to which the movement of individuals results in a change in state population. If economic conditions in one region are improving relative to condi-

[2]For papers on the decline in U.S. labor mobility based on the CPS data see Amior (forthcoming) and Frey (2020).

[3]For more detailed discussions of the differences between data sources see Meyer et al. 2015, Masnick (2013), Kaplan and Schulhofer-Wohl (2012) and Molloy et al. (2011).

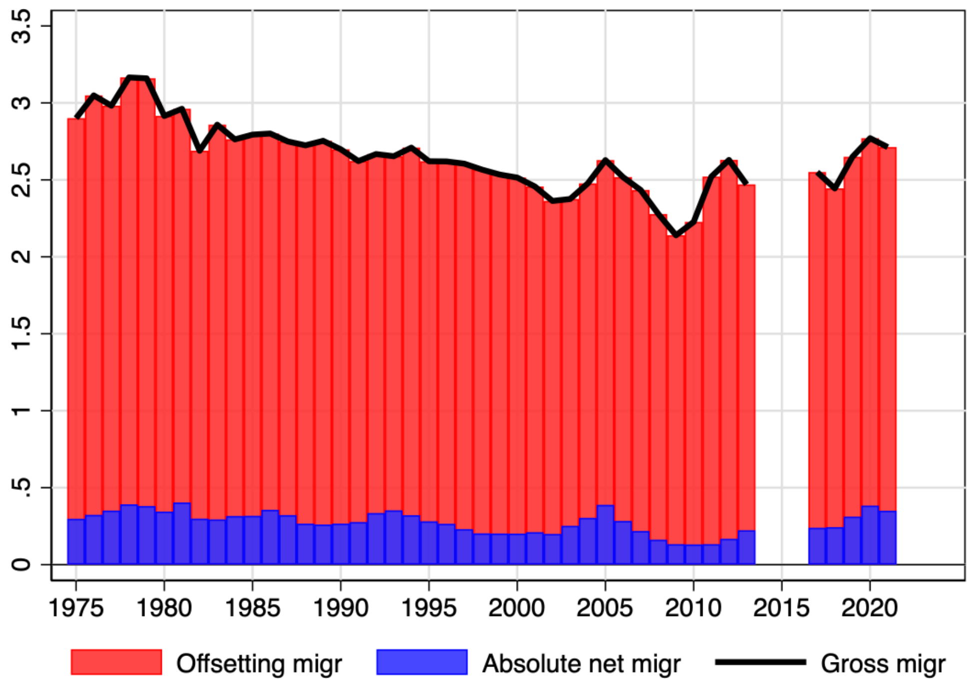

Figure 2: Gross, net and offsetting migration flows

Note: The figure plots the gross migration rate, the offsetting migration rate and the absolute net migration rate across states. Figures are expressed in percent of the US population. Data source is the IRS.

tions in anther region, it is natural to anticipate that there would be a net inflow of workers to the better-performing region. In contrast, Offsetting Migration is a measure of "churn" – worker migration flows that are matched by offsetting flows in the opposite direction.

Figure 2 plots all three rates – gross migration and its decomposition into the contributions of absolute net migration and offsetting migration – since 1975 using the IRS data. We see again the slight decline in the gross migration rate, matched by a similar trend in the Offsetting Migration Rate. We also see that the majority of gross migration is given by flows that offset each other, and only a small fraction of flows actually induce a change in a given state's population. On average, only 10-12 out of 100 moves reflect a net flow from one location to another. Net flows are a small fraction of overall labor migration and have remained fairly constant. In a theme that we will pick up again in section 3, this suggests that while overall labor mobility may be declining somewhat, the flows that help stabilize differences across states have not significantly changed. These data also suggest that models that rely on relative differences between locations as a driver for migration will be successful in explaining only a fraction of migration flows.

2.2 Can Demographic Factors Explain the Decline of Labor Mobility?

Gross migration rates in the IRS data fell by about 0.3 - 0.5 percentage points over the last 50 years. Possible explanations for this decline could be that the US population is aging, that more people now own their own homes – raising the cost of relocating – or the share of foreign-born individuals has changed. Figure 3 plots migration rates by age, by educational attainment, by housing status and by nativity. We report these rates based on CPS and ACS data. Blue lines reflect CPS data; green lines ACS data. While there are differences in the migration rates based on CPS data relative to ACS data, a number of patterns emerge regardless of the source of the data. First, the longer sample of CPS data show a decline in migration rates across demographic groups, mirroring the aggregate decline in migration rates we saw in Figure 1. Second, both ACS and CPS migration rates have generally leveled off starting in the early 2000s. Third, the data suggest that migration rates are lower for individuals that are older, are less-educated, own their own home, or were born abroad.

Figure 4 provides more detail for the various demographic categories based on the ACS data in the 2020-2024 period. Migration rates for various demographic groups are expressed as differences from the population-wide gross migration rate. We observe that younger adults are substantially more mobile than the average. Migration rates of those in the 21-25 age bracket exceed the average rate by 3 percentage points. Migration rates quickly fall as people become older, with people in their late 30s having migration rates that are representative of those of the average population. More-educated people are more mobile, with mobility rates of the college-educated exceeding those with less-than high-school education by more than 1.5 percentage points. Renters also have higher mobility rates than homeowners (by almost 2.5 percentage points). The difference between foreign-born and natives is substantially smaller, with natives having slightly higher migration rates.

To see the impact of demographic factors on the gross migration rate, we calculate counterfactual U.S. migration rates that fix the group-specific migration rates to those observed in the ACS in the 2020-24 period, but let the population share fluctuate as observed in U.S. time series going back to the 1970s. This exercise allows us to gauge to what extent demographic changes alone, rather than changes in migration rates across demographic groups, can account for the downward trend in migration rates. Figure 5 displays the counterfactual gross migration rates together with the actual migration rate as observed in the IRS data.[4] The 1976-1979 average for each migration

[4]Time series for population shares by demographic groups are from (i) the Survey of Epidemiology and End Results

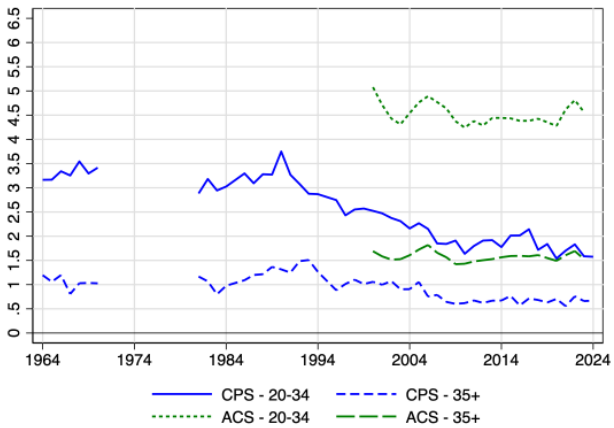

(a) Age

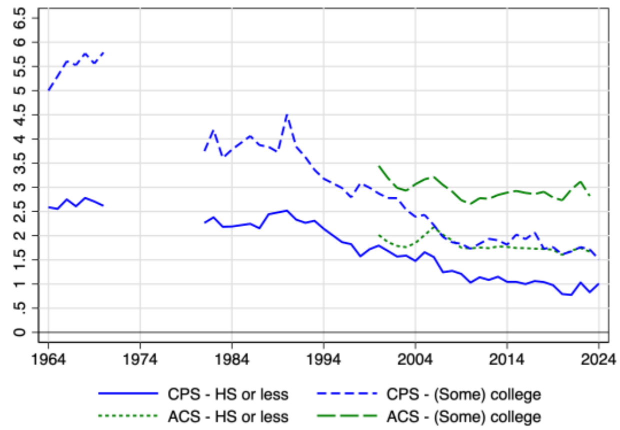

(b) Education

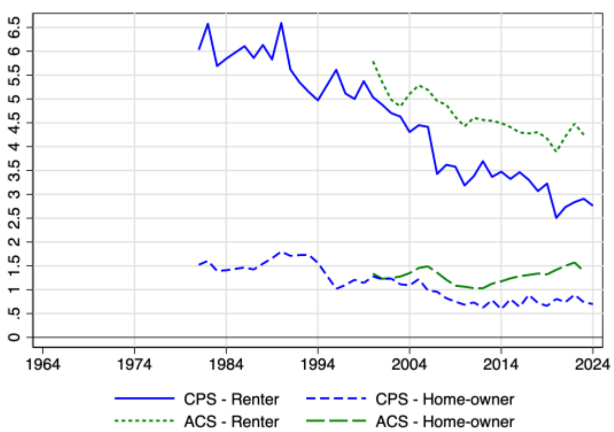

(c) Renters

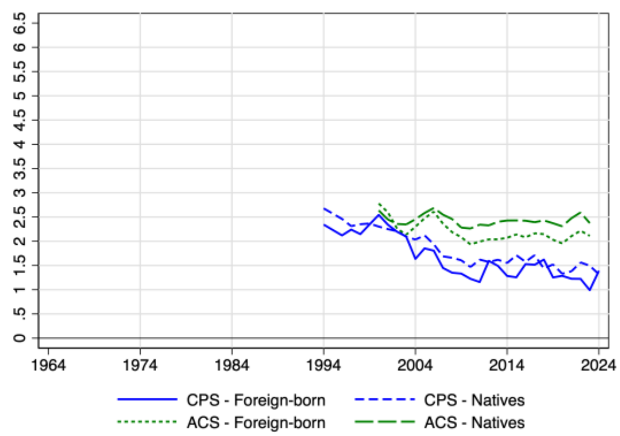

(d) Nativity

Figure 3: Migration rates by demographic group

Note: The figure plots gross migration rate for various demographic groups. Data sources are the CPS and the ACS.

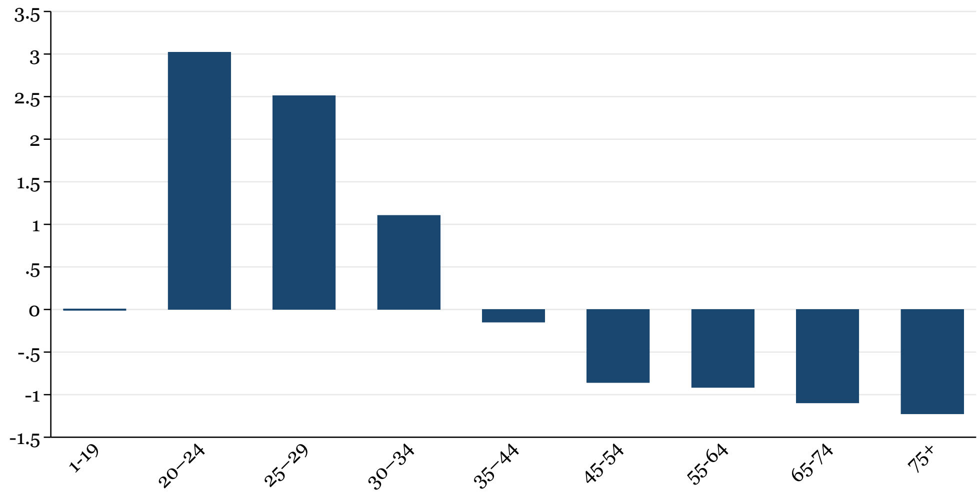

(a) Age

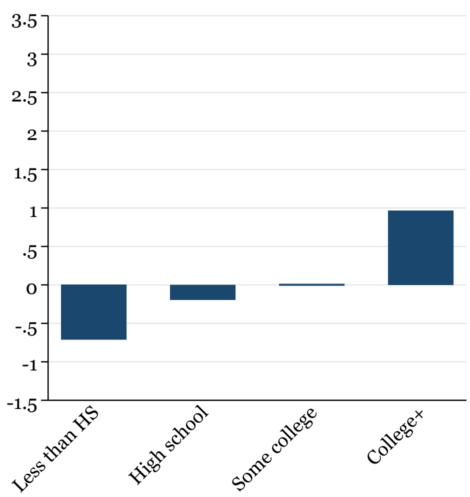

(b) Education

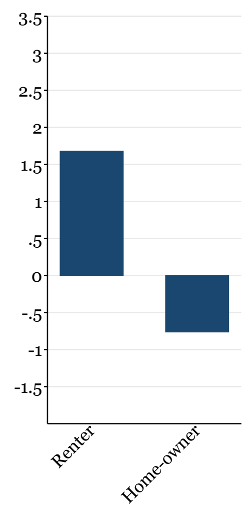

(c) Renters

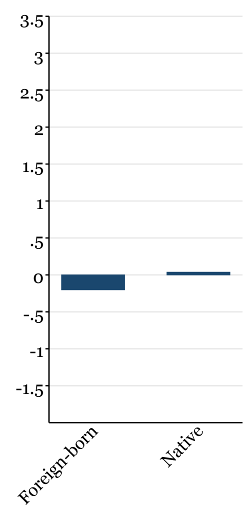

(d) Nativity

Figure 4: Migration rates relative to overall migration rates

Note: The figure plots gross migration rate for various demographic groups, pooled over the years 2020 - 2024. Migration rates are expressed as a difference from the overall migration rate. Data source is the ACS.

rate is normalized to zero to better visualize their change over the last 45 years. The blue dashed line shows the migration rate if only the age composition of the population had changed. The aging of the society implies a fall in the migration rate since the early 1980s because older people have lower migration rates. Aging by itself accounts for a drop of about 0.2 percentage points. However, this effect is more than counterbalanced by the impact of a more-and-more educated population. The share of people with no high school degree fell from 40 percent to 14 percent over the last 45 years, whereas the share with a college degree increased by 19 percentage points to 35 percent. Since college-educated people are more mobile, this should have raised the migration rate by about 0.35 percentage points since the mid 70s. Changes in the share of renters have had a minor impact on migration rates. Similarly, despite an increase of the foreign-born in the U.S. population since 1970 (from about 5 to 15 percent of the population), the impact on aggregate migration rates is minimal because foreign-born and natives have very similar migration rates.

Taking all of these demographic changes into account, the predicted change in the migration rate over the last 45 years is close to zero. To answer the question posed at the top of the section, Can demographic factors explain the decline in labor mobility? , our answer is no.

3 The Elasticity of Net Migration to Regional Shocks

The fact that gross migration rates have declined does not necessarily mean that labor has become less sensitive to economic conditions. In this section we examine the response of state-level net migration to plausibly exogenous changes in local labor demand. Our analysis will allow us to estimate the dynamics of net migration in response to the shock as well as the importance of net migration relative to the change in employment.

3.1 Estimation Method and Baseline Results

Following Foschi et al (2025), for any given left-hand-side variable Y_i, we estimate a set of horizon-specific regressions (i.e., Jordà (2005) local projections),

$$Y_{i\,t\quad h} = \quad {}^{Y}_{i\,h} + \quad {}^{Y}_{t\,h} + \beta^Y_h Z_{i\,t} + \delta^Y_h Z_{i\,t\quad 1} + \varepsilon^Y_{i\,t\quad h}. \tag{1}$$

(SEER) of the Cancer Institute (which in turn bases itself on U.S. Census population estimates) for age, (ii) the CPS for education, (iii) the housing ownership and vacancies survey by the U.S. Census for homeownership rates, and (iv) the American Community Survey (since 2000) and the decennial census (for 1970, 1980, and 1990, with linear interpolations) for the foreign-born.

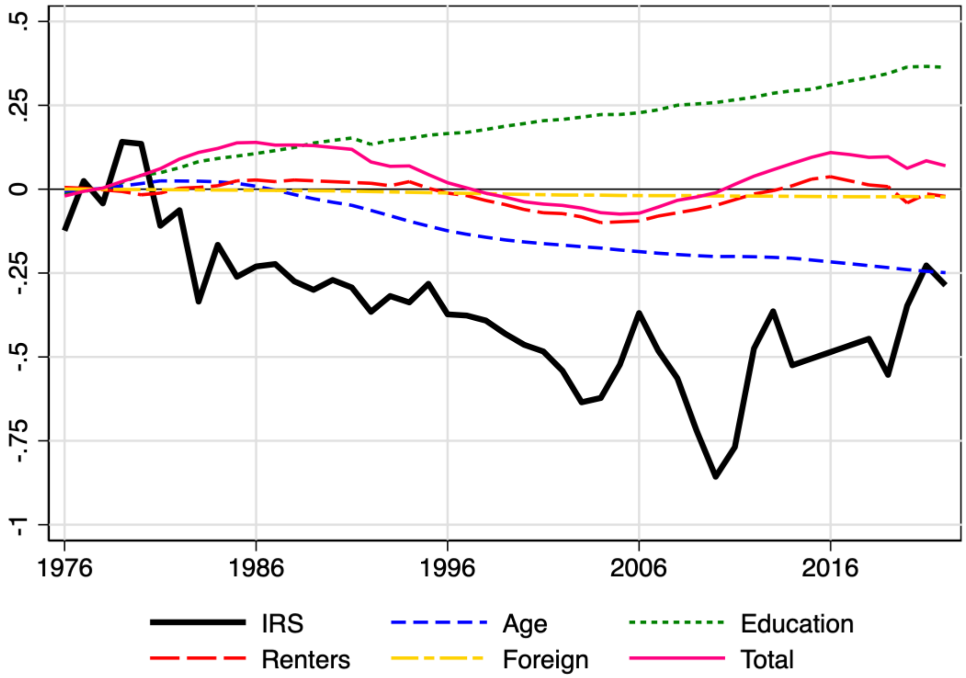

Figure 5: Counterfactual migration rates: Deviations from 1976-1979.

Note: The figure displays the gross migration rate from the IRS data together with a set of counterfactual migration rates. The counterfactual rates are calculated by fixing the migration rate for each demographic group estimated from the 2020-2024 ACS sample and then aggregating across groups using observed population shares back to 1976. The 'Total' counterfactual is the sum of all the other counterfactual migration rates. All rates are expressed in difference to the average rate over 1976 - 1979.

Here, i denotes a state, and h is a horizon in years. The key variable in this regression is the labor demand shifter $Z_{i\,t}$. For this analysis, we use the traditional industry-composition instrument suggested by Bartik (1993). The Bartik instrument has been widely used as a proxy for local labor demand shocks (see, among many others, Bound and Holzer 2000; Beaudry et al. 2014; Notowidigdo 2020; Dao et al. 2017; Amior and Manning 2018). The industry-composition instrument is obtained by constructing the expansion of employment that would have occurred in state i given the historical industrial composition of employment in that state if each industry grew at the national growth rate. For example, an increase in U.S. demand for wine will increase national demand for workers in vineyards and wineries, but it will have a disproportionate impact on employment in California, where most of the U.S. wine industry is located. More specifically, the

Bartik instrument is defined as

$$Z_{i\,t} = \sum_{j=1}^{J} s^j_{i\,t} \frac{\mathcal{E}^j_{t \quad i}}{\mathcal{E}^j_{t \quad 1 \quad i}}, \tag{2}$$

where $\mathcal{E}^j_{t \quad i}$ denotes national employment in industry j at time t, excluding state i, and $s^j_{i\,t}$ is the average employment share of industry j in state i over the preceding four years. Intuitively, the Bartik shock reflects the change in employment that state i would have experienced if its industry structure had grown at national (excluding i) industry-specific rates. In the following section we will delve more into the properties of the instrument.

The econometric specification in (1) includes a lag of the instrument as well as horizon-specific time and region fixed effects. Standard errors are clustered at the state level and the regressions are weighted by a region's population share relative to the national population. The data sources for each of the variables and a full description of the method used to create the Bartik instrument are included in the appendix.

Figure 6 shows the estimated local projections for four left-hand-side variables to a one unit increase in the Bartik industry composition instrument. That is, the figure plots the $\hat{\beta}_h$ coefficients from equation (1). The variables are the log change in the employment rate ($\ln E_{i\,t \quad h}$), the log change in the unemployment rate ($\ln(1 \quad ur_{i\,t \quad h})$), the log change in the labor force participation rate ($\ln LFP_{i\,t \quad h}$) and the log change in the population ($\ln POP_{i\,t \quad h}$). For now, we interpret the change in population as net migration, assuming that births and deaths are not directly affected by cyclical labor market shocks. [5] We later provide some evidence for this when we decompose the change in population into its individual components. Because, for each state i and each year t,

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

the estimated coefficients in the upper left panel will equal the sum of the estimated coefficients in the other three panels.

The plot in the upper left shows that the estimated employment response is slightly less than 1 percent on impact ($h = 0$) and then continues to rise over time.[6] The predicted change in state

[5] In the previous section on gross migration, we needed IRS/ACS data to capture both those leaving a state and those migrating into a state. To track net migration, we only require the change in each state's population and can rely on Census data.

[6] Because our Bartik instrument is based on a measure of jobs, the predicted change in jobs equals one, and the

employment is roughly 2.5 percent five years after the initial positive innovation in the instrument. The plot on the upper right shows that unemployment falls on impact and remains almost 1 percent below its initial value at all horizons. The bottom left panel shows that there is only a modest change in labor force participation.

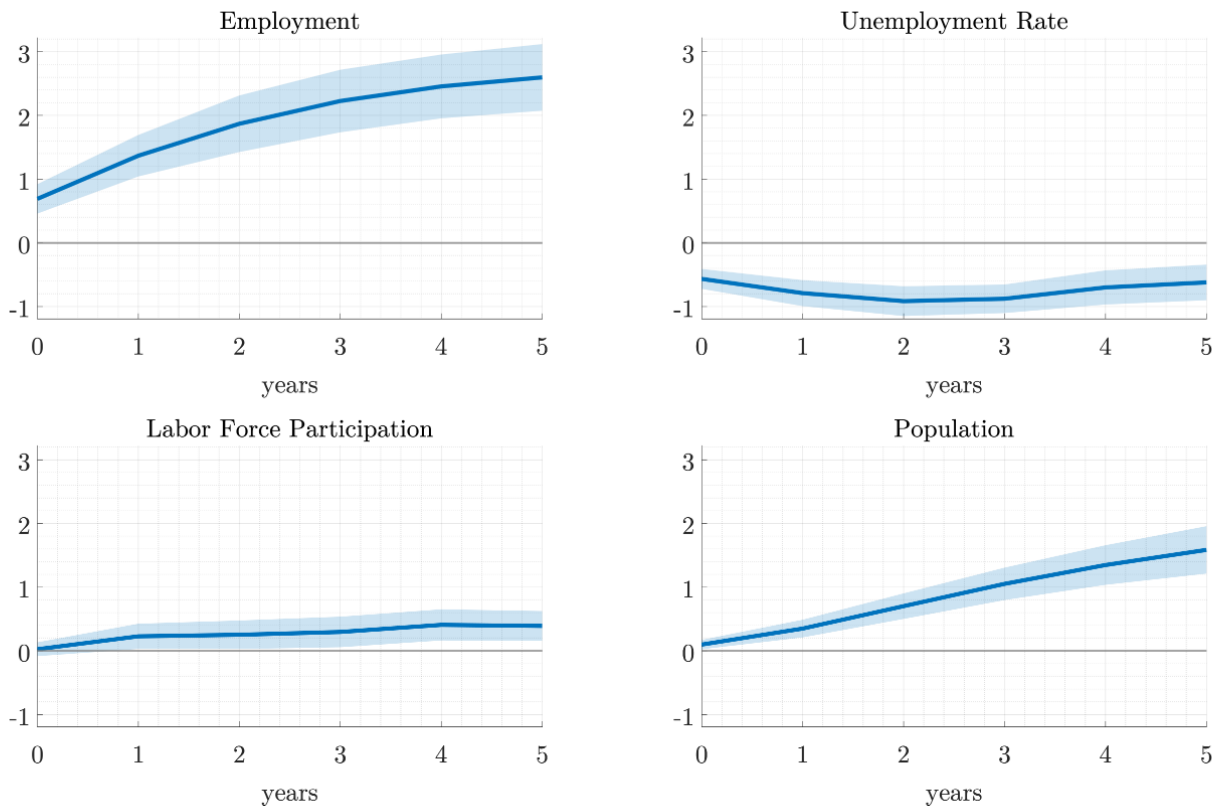

Figure 6: RESPONSE TO BARTIK INSTRUMENT: LABOR MARKET VARIABLES

Note: The figure plots the estimated $_h$ coefficients obtained from running (1) for each labor market variable at different horizons h (x-axis). The sample period is from 1976 to 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

Our main coefficient of interest is the migration coefficient β_h^{POP} plotted in the lower right panel. These estimates show how net migration in a state reacts to a one percent increase in the Bartik instrument. 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.[7]

For each variable other than employment, we can calculate the implied coefficients $\gamma_h^Y = \beta_h^Y / \beta_h^E$. By construction, $\hat{\gamma}_h^{POP} + \hat{\gamma}^{LFP} + \hat{\gamma}_h^{ur} = 1$. These coefficients have a natural interpretation

predicted change in employment will be less than 1.

[7]For a full discussion of the Bartik instrument and a comparison of our results to other studies, see Foschi et al. (forthcoming)

as the fraction of the overall employment response attributed to each of the three components. Figure 7 plots the estimates for the γ coefficients associated with the β coefficients in Figure 6. [8] The shaded regions reflect one standard deviation error bands.[9]

On impact, most of the increase in employment demand is accounted for by a reduction in state unemployment. In-migration and increased labor participation account for the remaining 20 percent. After five years, more than 60 percent of the increase in employment demand is met by a net increase in population, with a much smaller fraction due to a fall in unemployment.

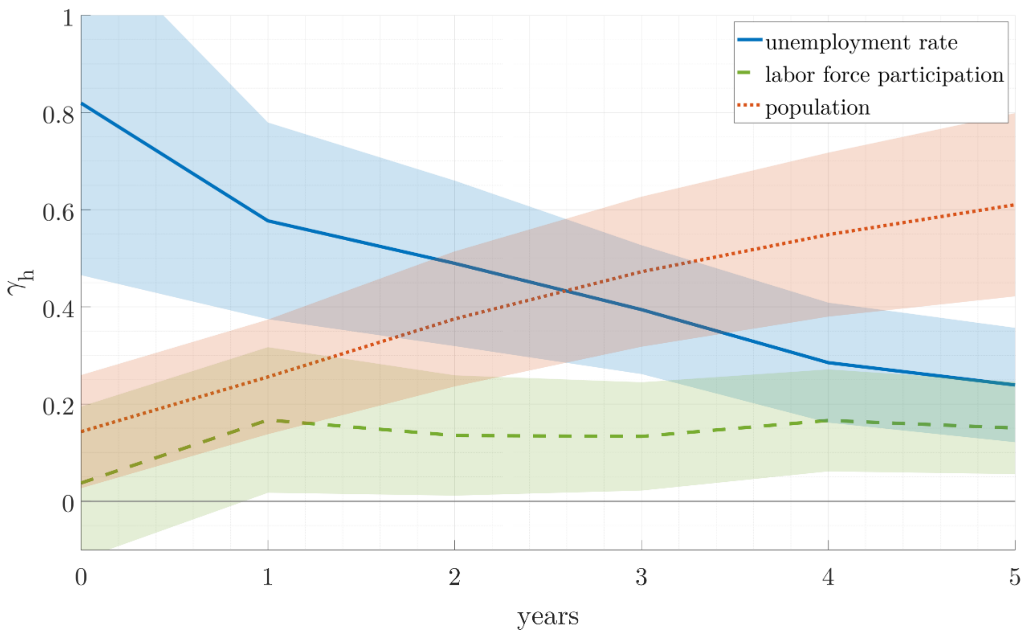

Figure 7: Ratio between each labor market response and the employment response

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

A central question is the extent to which labor mobility has changed over time and what that might imply for macroeconomic adjustment. We have already established that there is a modest decline in the rate of gross migration. But, how much has net migration, the component that arguably responds to differences between locations, changed? To get at this question, we consider a long sequence of rolling regressions. To produce the estimates for the longer time period, we are forced to use 'jobs' rather than 'employment' as our measure of labor demand. Figure 8 shows a

[8] These are identical to the results in Foschi et al. (forthcoming).

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

sequence of estimates of the migration elasticity over a longer time period. Each point in the figure is a point estimate for the 5-year population response share (i.e., $\hat{\gamma}_5^{POP}$) over a 10-year sub-sample of data. The lines represent the 90 percent confidence intervals for each estimate. The estimated net migration elasticity has remained remarkably consistent over the past 80 years and shows no indication of a decline in sensitivity.[10] Looking at the behavior of $\hat{\gamma}_5^{POP}$ around recessions, it also seems to have no apparent relationship with the business cycle.[11]

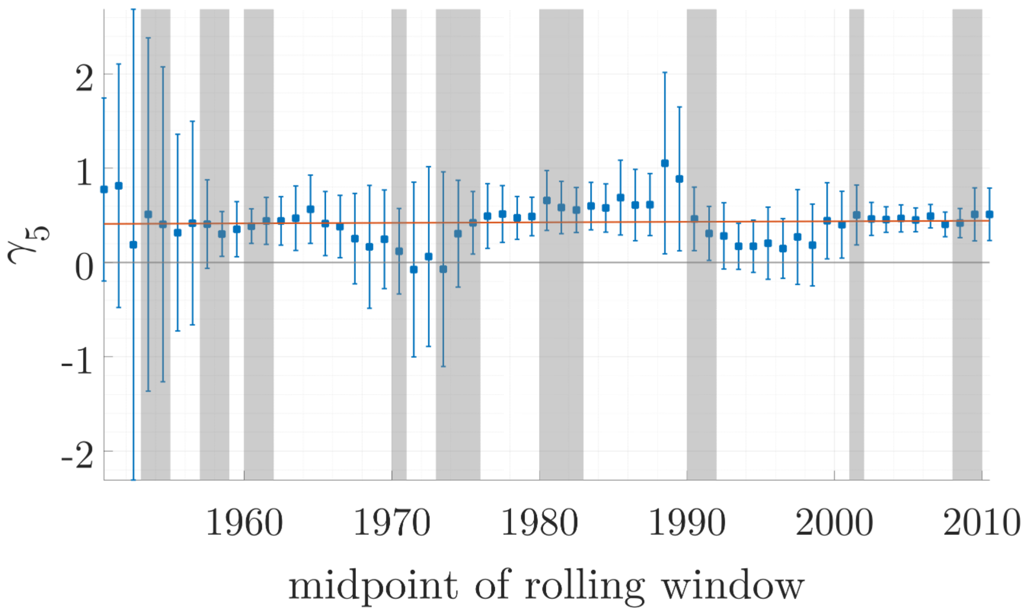

Figure 8: Ratio of population and jobs responses (γ_5^{POP})

Note: The figure plots estimates for $\hat{\gamma}_5^{POP}$ for 10-year data windows. Each point is centered at the mid-point of the 10-year window on the horizontal axis. Bars represent 90% confidence intervals. The red line is a weighted least square fit, where the weights are the inverse of the standard errors. Shaded areas indicate NBER recessions.

3.2 Other Labor Demand Shocks

The Bartik instrument may only capture one type of shock to labor demand and may miss other shocks that affect particular regions in the United States at particular points in time. In previous work, we estimated the population response at different levels of aggregation and in response to different types of shocks. We summarize those results in Table 1 below. The table shows the beta

[10]The estimates are clustered around $\hat{\gamma}_5^{POP} \approx 0.4$ rather than 0.6. This is because we use jobs rather than employment in the regressions. While net migration accounts for 60 percent of the increase in employment, the jobs to employment ratio also rises following a Bartik shock and so net migration accounts for only 40 percent of the increase in jobs.

[11]We also confirm this by computing lead-lag correlations between $\hat{\gamma}_5^{POP}$ and national GDP growth and we find no evidence of cyclicality.

coefficients on employment, unemployment and the change in population, and the far rightmost column shows the implied gamma elasticity. All of the results are for a time horizon of five years. Part A of the table shows results at the state level, the commuting zone (CZ) level and the county level. There is some variation in the magnitude of the employment response to the shock, but it is always positive; the unemployment response is small and negative, and the population response if relatively large and positive. The gamma coefficients are in the range of 0.54 to 0.61.

Table 1: Long-run labor market responses to different shocks and at different levels of aggregation

	Empl	5 ur	Pop	γ_5^{POP}
A. Different levels of aggregation				
State	2.60 (0.32)	0.62 (0.17)	1.59 (0.23)	0.61 (0.11)
CZ	1.73 (0.17)	0.15 (0.07)	0.99 (0.09)	0.57 (0.08)
County	1.15 (0.08)	0.07 (0.03)	0.62 (0.05)	0.54 (0.06)
B. Alternative shocks				
Defense spending (State)	2.00 (0.71)	0.50 (0.24)	1.35 (0.46)	0.68 (0.33)
Defense spending (CBSA)	0.89 (0.59)	0.44 (0.38)	0.04 (0.24)	0.05 (0.28)
Housing net worth (County)	0.47 (0.13)	0.18 (0.04)	0.29 (0.05)	0.62 (0.20)
Import competition (CZ)	2.41 (1.08)	0.03 (0.23)	1.25 (0.59)	0.52 (0.34)

Note: The first three columns of the table display the estimates of the $_h$ coefficients obtained by running (1) for each of the dependent variables listed at the top for horizon $h = 5$. The fourth column is the γ coefficients for net migration defined as the ratio between the ^{POP} and the employment ^{Empl}. The sample period is from 1976 till 2016. Standard errors (in parentheses) for the ’s and γ are clustered at the level of the region. All regressions are weighted by population. For the housing net worth shock and the import competition shock, the responses are measured relative to the year in which the shock is thought to have started. Specifically, for the housing-net-worth shock, the response is measured in 2012 (5 years from 2007); for the import-competition shock, the response is measured in 2006 (5 years from 2001).

Part B of the table shows the results based on alternative shocks that have been used in the literature as plausibly exogenous shocks to local demand. These include a shock to defense spending at the state level (Nakamura and Steinsson 2014) and at the core-based statistical area (CBSA) level (Auerbach et al. 2020); a shock to housing net worth at the county level (Mian and Sufi 2014 and Bhattarai et al. 2021, instrumented with the Saiz 2010 instrument); and the China import competition shock at the commuting zone level (Autor et al. 2013, 2021). These shocks are somewhat different as they occur at particular points in time, and therefore one has to be concerned with pre-trends and other controls. For a full discussion of the regression specifications, please see Foschi et al. (forthcoming). Here we focus on the main result, which is that despite differences in the specification of the shock, the elasticity of net migration is again in the range of 0.52 to 0.68 (with the exception of the low migration response to the defense shock at the CBSA level).

We take this as validation that, despite the gradual decline in the numbers of people moving across state lines over time as a share of the population, the movement of the population in response to differences in economic conditions between locations has not diminished.

3.3 The Response of Wages, Rents, and House Prices

If Bartik shocks do indeed reflect regional shocks to labor demand then one would expect other telltale signs in affected regions. Among other things, one would naturally expect a positive innovation in labor demand to predict an associated increase in wages. To the extent that improved labor market conditions encourage the in-migration of workers, we should also see higher rental rates and higher house prices. Figure 9 plots local projection impulse responses for the log change in nominal wages, house prices and rental rates.[12] For wages, we calculate the ratio of total wages and salaries to the number of jobs – both series come from the BEA; the house price data are from the Federal Housing Finance Agency (FHFA); for the rent data we use fair-market rents from the Department of Housing and Urban Development (HUD). For wages and house prices, the sample runs from 1976-2016; for rents, the sample is somewhat shorter from 1983-2016.

All three series react in the "correct" direction. The wage estimates suggest higher wages of roughly 1 percent on impact. The wage path peaks after roughly 4 years and then reverts to its previous level.[13] Rental prices rise by perhaps as much as half a percent on impact and peak after 4 years. There is some suggestion that rental rates remain above trend indefinitely.

House prices also move in the expected direction but the dynamics are puzzling from the

[12]These are nominal variables but since we include time fixed effects, we are essentially adjusting for a common national inflation term.

[13]Properly adjusting aggregate wages for composition bias (e.g., Barsky et al. 1994), and remittance-smoothing (e.g., Beaudry and DiNardo 1991; Kudlyak 2024) is a perpetual source of frustration for labor economics. See Basu and House (2016) for additional discussion.

Figure 9: Wages, House Prices, and Rent

Note: The figure plots estimated $_h$ coefficients obtained from running (1) for wages, house prices, and rent prices for different horizons h (x-axis). The sample period is 1976-2016 (1983-2016 for rent). Shaded areas represent 90% confidence intervals.

standpoint of neoclassical economics. Since one would expect house prices to capitalize the value of future expected rents, it would be more natural to observe a once-and-for-all increase at $h = 0$ (i.e., a random walk) by the average increase in the rent plot – say by roughly 1 percent. This doesn't happen in our data. House prices rise gradually and increase after four years by nearly 2-3 percent. This could be caused by time-averaging or selection in the house price data.

On the whole, these estimates suggest that the Bartik shocks are functioning as we would expect – an innovation implies elevated wages and higher house prices and rental prices as workers arrive seeking employment.

3.4 The Elasticity of International Migration to Regional Shocks

In the previous section, we interpreted the change in population as net migration. Starting in 1991, the U.S. Census provides a more detailed breakdown of state-level population changes into births, deaths, net domestic migration, net international migration and a (statistical) residual. This allows us to decompose the population change into its components for a shorter sample period. The main data source for net domestic migration are the mailing address information from the IRS tax return data for ages 0 to 64 and medicare enrollment data for individuals ages 65 and older. Net international migration consists of both foreign-born and native migration to and from the United States, including undocumented immigration as well as the net movement of the Armed Forces population.

The growth in a state's population between any two periods is the sum of domestic net migration (i.e., inflows minus outflows of U.S. citizens), international net migration, and natural population change (i.e., births minus deaths). Figure 10 shows the estimated local projections to the Bartik shock for population growth and the three sub-groups, each measured as a fraction of the state's population. The figure shows that, following a positive innovation to the Bartik instrument, population gradually rises by just under 1 percent after 5 years. The vast majority of the change in population comes from domestic net migration – the movement of U.S. residents. International net migration accounts for roughly only 0.1 percent and natural population change is essentially zero (i.e., births and deaths don't appear to react to the Bartik shock).[14]

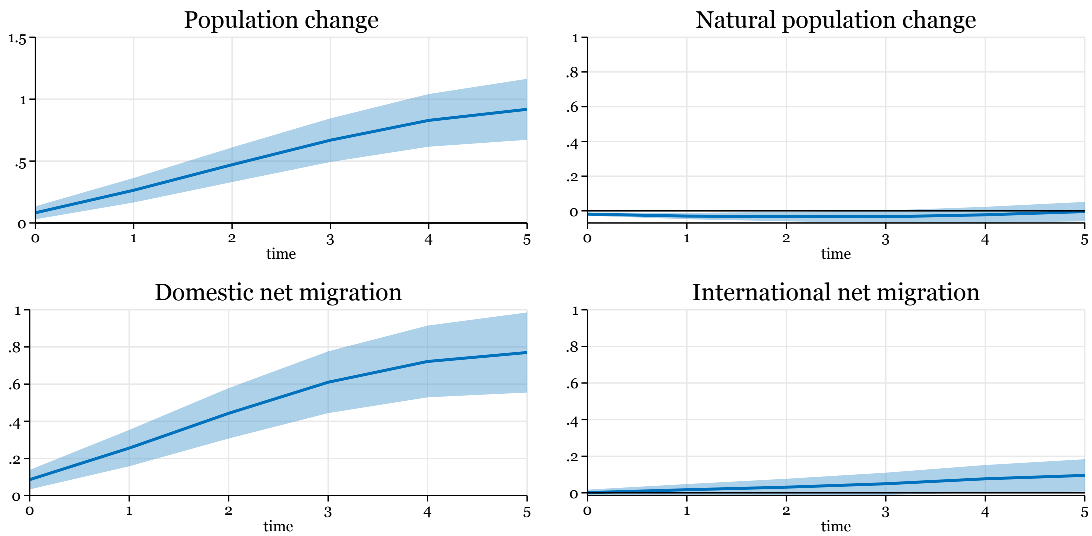

Figure 10: Response to Bartik instrument: Components of population change

Note: The figure plots the estimated $_h$ coefficients obtained from running (1) for the change in population and its components at different horizons h (x-axis). The sample period is from 1990 till 2016, with projections going up to 2021 (for $h = 5$). Shaded areas represent 90% confidence intervals.

3.5 The Migration Elasticity of Demographic Groups

Figure 11 shows estimates of local projection responses for the four demographic subpopulations we considered in Section 2.2. For the age population groups, we use state-level data from the Surveillance, Epidemiology, and End Results (SEER) database from the National Cancer Institute. This sample allows us to use data back to 1976. For renters and owners, foreign-born and domestic born, and education groups we need to rely on ACS data. This forces us to restrict the sample to the post-2000 period. This makes the estimates for these series somewhat noisier than those for

[14]The remaining share of population change is the statistical residual.

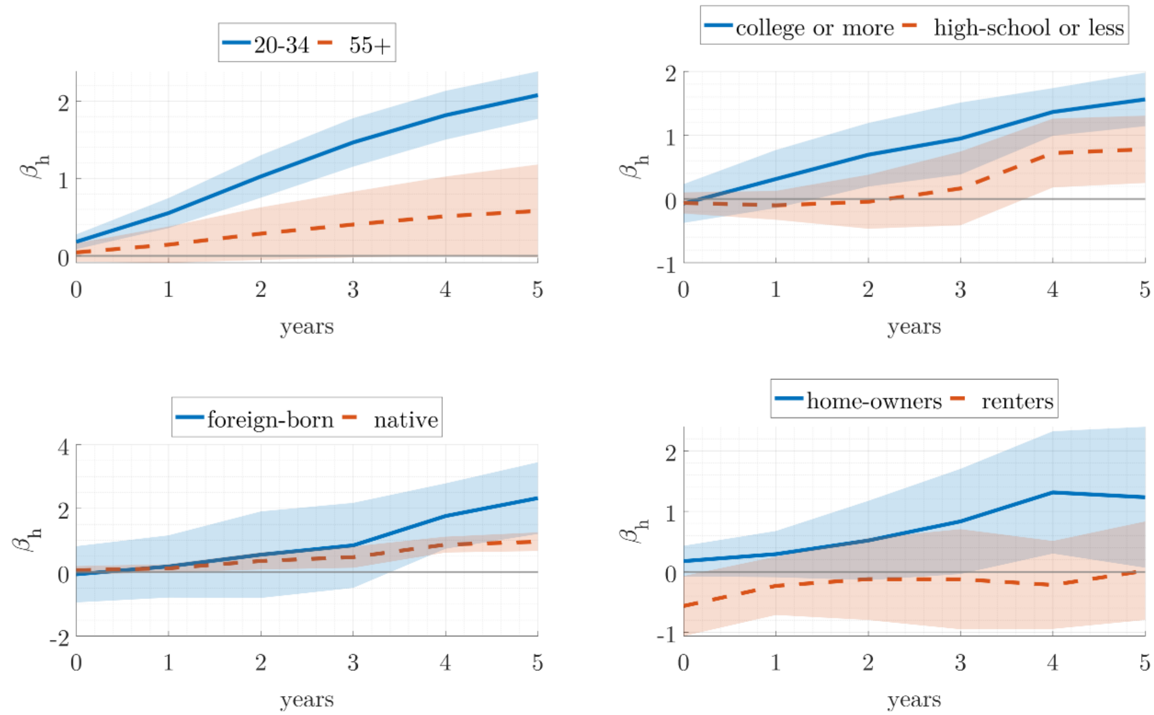

Figure 11: Different demographic groups

Note: The figure plots estimated ${}_h$ coefficients obtained from running (1) for several demographic groups for different horizons h (x-axis). The sample period is 1976-2016 for age groups and 2000-2016 for the other groups. Shaded areas represent 90% confidence intervals.

the longer samples back to the 1970's.[15]

The population response by age in the upper left panel shows, as we would expect, that younger workers react more to labor demand shocks than do older workers. Five years after the innovation, the population of workers aged 20-34 has increased by roughly 2 percent while the population of workers age 55 or older has increased by approximately one half percent. Foreign-born workers and college educated workers also seem to be more responsive though the estimates are less precise.[16]

Interestingly, we do not find evidence that renters move in response to labor demand shocks

[15] Due to the constraints of the ACS data, we also make additional adjustments to compute the estimates for education groups, native vs foreign-born, and renters vs home-owners. First, we exclude data from 2020 and 2021 because the ACS data are less reliable for the COVID period. Second, since the ACS data are only available starting in 2000, we modify our specification by replacing the state fixed effects in (1) with state-specific trends and controls for the growth of the dependent variable between 1980 and 2000 (which we compute from the decennial Census) and the average value of the Bartik instrument in the same period. This allows us to control for pre- and post-2000 trends in a way comparable to the fixed effects we use when relying on data from the 1970s.

[16] Amior (forthcoming) considers the impact of foreign immigration on U.S. employment back to 1960 and finds a minimal impact. Basso et al. (2019) and Basso and Peri (2020) also study the responsiveness of foreign-born workers to economic conditions. Cadena and Kovak (2016) find higher responsiveness of Mexican immigrants to local labor demand shocks than unskilled domestic workers.

more than home-owners (in fact the point estimates are reversed). Based on the stark differences in gross migration rates we observed in Section 2.2, we anticipated that renters would have a much higher elasticity. This is not what we observe in the impulse responses. Whether this result is an actual signal that renters react less to labor demand shocks or whether it is simply a fluke in the data is not clear at this time.

4 How Important are Bartik Shocks?

The Bartik shock captures state-specific shifts in labor demand and is strongly positively correlated with both employment and net migration, as previously shown. In this section, we assess its role in shaping trends in employment and migration over time.

As explained in the previous section, the Bartik instrument is constructed as

$$Z_{i\,t} = \sum_{j=1}^{J} s^{j}_{i\,t} \frac{\mathcal{E}^{j}_{t\quad i}}{\mathcal{E}^{j}_{t\;\;1\;\;i}}.$$

Intuitively, the Bartik shock reflects the change in employment that state i would have experienced if its industry structure (represented by $s^{j}_{i\,t}$) had grown at national (excluding i) industry-specific rates (represented by $\mathcal{E}^{j}_{t\quad i}/\mathcal{E}^{j}_{t\;\;1\;\;i}$).

Figure 12 shows the cross-sectional standard deviation of the Bartik shock together with the cross-sectional standard deviation of state-level employment growth. Two insights emerge: First, the variation in predicted growth rates across states (i.e., the Bartik shock) is only one-third of the variation in actual growth. This suggests that differences in industry composition across states only capture a small fraction of the overall differences in observed employment growth across states. Second, both series exhibit a downward trend in dispersion over time, with cross-sectional variation in the latter part of the sample about half as large as in the earlier years.[17]

To quantify the explanatory power of the Bartik shock, we follow the approach of Gorodnichenko and Lee (2020) and use a two-step procedure to compute its contribution to the forecast error variance at horizon h. First, we estimate local projections for employment growth and population growth excluding the Bartik instrument. We calculate the forecast error $\hat{\epsilon}^{Y}_{i\,t\;\;h}$. We then

[17] This might be related to the shift from manufacturing to services; however, the downward trend remains even when using only manufacturing industries to construct the Bartik, suggesting the transition to services might not be the whole story.

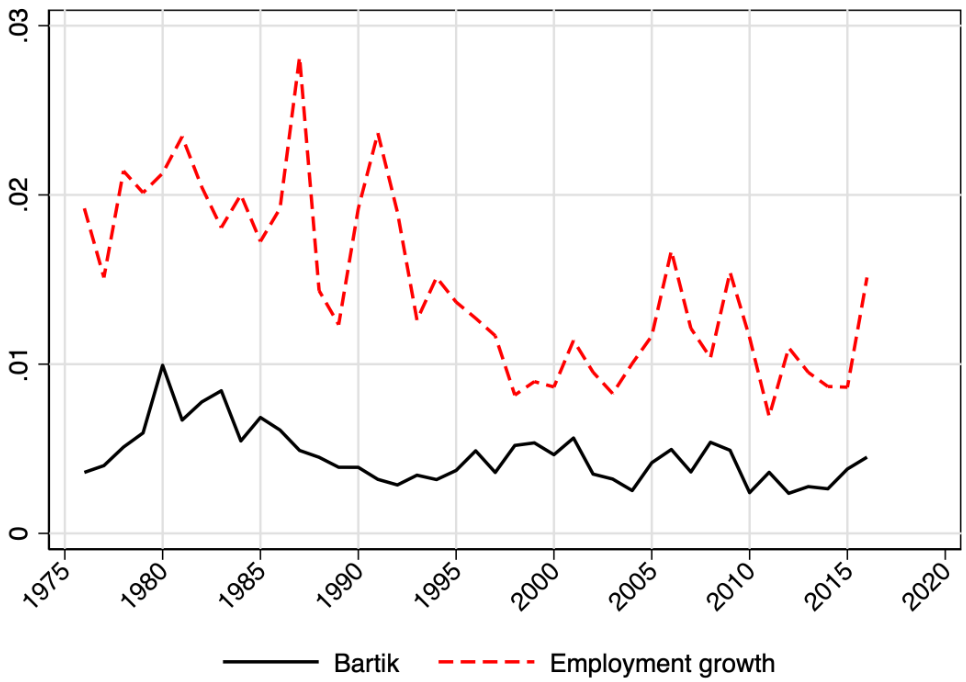

Figure 12: Cross-sectional standard deviation

Note: The figure displays the cross-sectional standard deviation of the Bartik instrument and employment growth.

regress this error on the current and lagged values of the instrument:

$$\hat{\epsilon}^{Y}_{i\,t\quad h} = \psi_0 Z_{i\,t\quad h} + ... + \psi_h Z_{i\,t} + \nu_{i\,t\quad h}.$$

Table 2 reports the partial R^2 from this regression. For employment growth, the Bartik shock explains approximately 38% of the forecast error variance at horizons 2 and beyond. In contrast, it explains a more modest 17% of the variance in population growth. These results underscore the Bartik instrument's strength in capturing differences in labor demand that drive regional employment dynamics, and to a lesser extent, migration flows.

4.1 Drivers of the fall in the Bartik shocks

What explains the decline in cross-sectional dispersion of the Bartik shocks shown in Figure 12? It could be that industrial composition across states has become more similar or it could be that industries are growing at more similar rates. To get at this question, we use variance decomposition to compare the variance in the early period with the later period. For simplicity, we compare the early sample period (from 1977 - 1986) that was characterized by high dispersion (a cross-

Table 2: Forecast error variance decomposition

Horizon	Employment	Net migration
0	0.057	0.109
1	0.298	0.171
2	0.366	0.175
3	0.388	0.176
4	0.388	0.169
5	0.384	0.164

Note: The table displays the forecast error variance decomposition at horizons $h = 0, 1, ..., 5$ and for employment growth and population growth as left-hand-side variables.

sectional standard deviation of about 0.0068) and the late sample period (from 2007 - 2016) that was characterized by low dispersion (0.0037).

We calculate the cross-sectional variance of the Bartik shock in the early period as

$$V_1 = \frac{1}{T} \sum_{t=1977}^{1986} Var_t(Z_{i\,t}),$$

and for the later period as

$$V_2 = \frac{1}{T} \sum_{t=2007}^{2016} Var_t(Z_{i\,t}),$$

where the subscript 1 or 2 indicates whether this is the early or late period. We can then define two counterfactual variances that use either early-period shares with late-period shifts or late-period shares with early-period shifts:

$$V_{s1g2} = \frac{1}{T} \sum_{t=2007}^{2016} Var_t \quad \sum_{j=1}^{J} s^j_{i\,t\quad 30} \frac{\mathcal{E}^j_{t\quad i}}{\mathcal{E}^j_{t\quad 1\quad i}} \Bigg)$$
$$V_{s2g1} = \frac{1}{T} \sum_{t=2007}^{2016} Var_t \quad \sum_{j=1}^{J} s^j_{i\,t} \frac{\mathcal{E}^j_{t\quad 30\quad i}}{\mathcal{E}^j_{t\quad 31\quad i}} \Bigg).$$

Then, the total change in the variance can be decomposed into two components:

$$V_2 \quad V_1 = \underbrace{\frac{1}{2}\left((V_{s2g1} \quad V_1) + (V_2 \quad V_{s1g2})\right)}_{\text{constant shift}} + \underbrace{\frac{1}{2}\left((V_{s1g2} \quad V_1) + (V_2 \quad V_{s2g1})\right)}_{\text{constant share}}. \tag{3}$$

The first term of the composition reflects the change in the variance that would have been observed if only the shares had changed, either evaluated at early-period or later-period shifts. The

Table 3: Variance decomposition of the Bartik

V_1	0.466	
V_2	0.140	
V_{s1g2}	0.148	
V_{s2g1}	0.178	
$V_2 \quad V_1$	-0.326	
constant shift component	-0.148	45%
constant share component	-0.178	55%

Note: The table displays the variance decomposition of the Bartik instrument outlined in (3). For easier readability, the Bartik has been multiplied by 100 for these calculations, i.e. it is expressed as a percentage.

second term reflects changes in the variances that would have been observed if only the shifts had changed, either evaluated at early-period or later-period shares.

Calculating these terms suggests that, of the overall fall in the cross-sectional standard deviation from 0.0068 to 0.0037, about half can be attributed to changes in the shifts and half to changes in the shares, as shown in Table 3. Hence, industries are growing at more and more similar rates and states' industry composition has become similar.

4.2 The Bartik shocks and gross migration

In Section 3, we examined how net migration responds to regional labor demand shocks. As expected, positive shocks lead to net in-migration, while negative shocks result in net out-migration. Notably, this response is largely driven by changes in in-migration rather than out-migration: migrants are more likely to move toward states with favorable shocks than to actively leave states with negative ones. However, these patterns in net migration do not necessarily imply that a decline in the volatility of Bartik shocks—illustrated in Figure 12—leads to a decline in gross migration.

We now turn directly to the question of whether a state's gross migration rate is affected by changes in labor demand, irrespective of their direction. The underlying idea is that labor demand fluctuations, even when offsetting across states, generate turnover or "churn" in labor markets, prompting migration in both directions. [18] If labor demand becomes more spatially uniform over time, then this mechanism would predict a decline in gross migration.

To test this hypothesis, we estimate a variation of regression (1):

[18] See Azzopardi et al. (2020) for a survey of recent data on job flows and migration.

$$\frac{\sum_{s=0}^{h} \text{in-migrants}_{i\;t\;\;s} + \text{out-migrants}_{i\;t\;\;s}}{\text{population}_{i\;t\;\;1}} = \quad_{i\,h} + \beta_h |Z_{i\,t} \quad Z_t| + \delta_h |Z_{i\;t\;\;1} \quad Z_{t\;\;1}| + \varepsilon_{i\;t\;\;h}. \tag{4}$$

The dependent variable is cumulative gross migration (in- plus out-migration) in state i over the period t to $t + h$, scaled by the initial population in year $t \quad 1$. The key explanatory variable is the *absolute deviation* of the Bartik shock in state i from the national average in year t. We subtract the national mean because migration decisions are driven by relative—not absolute—labor demand conditions across states. We take the absolute value because we are interested in whether any change—positive or negative—in relative labor demand leads to more migration activity.

As in previous regressions, we include state fixed effects to absorb average, time-invariant differences in gross migration rates and relative labor demand changes across states. For example, smaller states often have higher gross migration rates due to their smaller size and larger shocks due to their less diversified industrial structures. Conversely, larger states tend to have more stable industry mixes, leading to Bartik shocks that closely track the national average—and lower gross migration rates due to their bigger size. Including state fixed effects ensures that our estimate of β_h is not biased by these cross-sectional relationships.

We do not include time fixed effects, however, because the time-series variation in the cross-sectional dispersion of Bartik shocks is central to our identification. In particular, in years where the national economy is more uneven, the value of $|Z_{i\,t} \quad Z_t|$ is mechanically higher across states. If such years are also associated with more migration, this variation contributes directly to our estimate of β_h, helping us identify the relationship between shock dispersion and gross migration.

Figure 13 displays the estimated impulse response. There is a clear positive response that grows over 8-9 years before leveling off. A 1 percent change in predicted labor demand raises gross migration by about 1.2 percent of population over 8-9 years.

This suggests that the fall in the volatility of regional shocks might have contributed to the observed decline in gross migration. We can quantify this effect through a back-of-the-envelope calculation. The mean of $|Z_{i\,t} \quad Z_t|$ declined from 0.005 to 0.002 from the early 80s till the 2010s, in line with the fall in the cross-sectional dispersion. This suggests a fall in the gross migration rate of about 0.36 ($= 100 \times 0.003 * 1.2$) percentage points, which is in the range of the actually observed fall in the gross migration rate as measured by the IRS data and depicted in Figure 1.

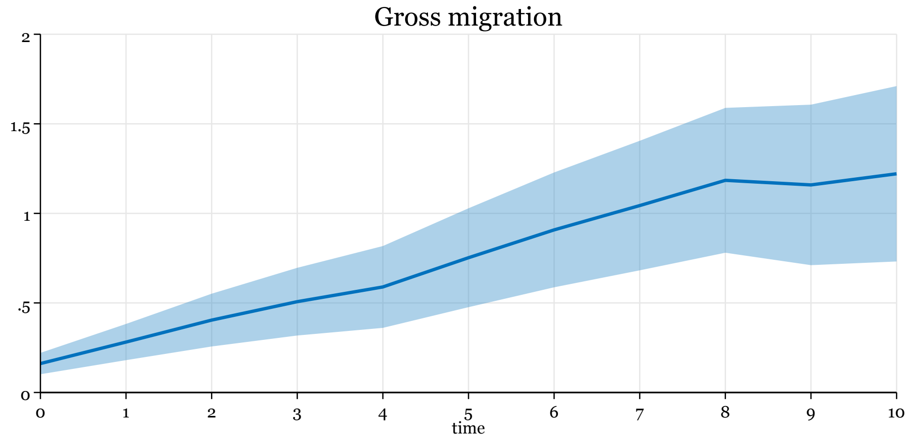

Figure 13: Response to the absolute value of the Bartik instrument

Note: The figure plots the estimated $_h$ coefficients from running regression (4).

Hence, the emergence of a more uniform business cycles is a promising candidate for explaining the fall in aggregate gross migration rates.

5 Conclusion

Many economists have raised concerns about the decline in gross migration rates, taking this as an indication of declining U.S. labor market dynamism. While it is true that U.S. gross migration rates have fallen over the past 50 years, the decline is relatively modest. Using the largest datasets with the most coverage, we find that gross migration has declined from roughly 3 percent in the late 1970s and early 1980s to roughly 2.5 percent today. Larger estimates of the decline, indicating a drop of 2-2.5 percent – nearly a 2/3 reduction in migration rates – are based on CPS data that is understood to be a less reliable measure of migration. The decline in gross migration is primarily associated with a decline in "offsetting migration" or churn, by which we mean migration flows from state a to state b that are matched by reverse flows from b to a which leave the state population levels unchanged. In contrast, *net* migration flows – flows that entail changes in state-level population – show little to any sign of falling.

The data also indicate that, while there may have been a modest decline in labor mobility in the U.S., the elasticity of net migration to identified labor demand shocks does not appear to have declined. Further, the estimates show that labor demand shocks have an impact on prices in the direction that one would expect. Following an increase in labor demand, house prices and rent both rise reflecting pressure to accommodate incoming workers. The robustness of net migration over time and across different kinds of shocks is reassuring, in that it suggests that workers are relocating in response to local shocks now to the same extent as they did in the past. Our analysis also suggests that the incentives to move may have declined over time. The cross-sectional dispersion in Bartik shocks has fallen. This is partly due to an increased similarity in industry composition across states, and partly due to shocks becoming more correlated across industries.

Our findings carry important implications for both spatial and labor market policy. Migration continues to serve as a critical mechanism for regional economic adjustment, and policymakers should not be too quick to conclude that labor mobility is a thing of the past. This does not mean, however, that policymakers should not be concerned about frictions that might prevent individuals from moving. Efforts to reduce barriers to migration – such as improving housing availability and affordability, addressing information frictions, and offering relocation support – could enhance the responsiveness of labor markets even more.

Moreover, these findings do not imply that place-based policies are unnecessary or that aggregate, nationwide interventions are sufficient. In the face of falling labor demand, our local projection estimates predict a substantial fall in population but they also show that local wages fall and local unemployment rates are slow to recover. Workers that can move – the more affluent, the higher educated, the young – can perhaps avoid the consequences of a negative regional shock. Workers that are less mobile – lower income, lower educated and older workers – may be trapped for years in communities hollowed out by falling labor demand. As such, place-based policies may still play an essential role in revitalizing contracting local labor markets by supporting resident populations and attracting new workers through tailored investments and employment opportunities.

As policymakers confront persistent regional disparities and economic shocks, it is crucial to acknowledge the enduring role of migration – as well as its limitations. By strengthening mobility where it is robust and eliminating barriers where it is constrained, economic and labor market policies can continue to sustain the dynamism of the American workforce.

References

M. Amior. The Contribution of Immigration to Local Labor Market Adjustment. *Journal of Labor Economics*, forthcoming.

M. Amior and A. Manning. The persistence of local joblessness. *American Economic Review*, 108 (7):1942–1970, 2018.

J. Angrist, P. Hull, and C. Walters. Methods for Measuring School Effectiveness. *Handbook of the Economics of Education*, 7:1–60, 2023.

J. D. Angrist and P. Hull. Instrumental Variables Methods Reconcile Intention-to-Screen Effects Across Pragmatic Cancer Screening Trials. *Proceedings of the National Academy of Sciences*, 120 (51):e2311556120, 2023.

A. J. Auerbach, Y. Gorodnichenko, and D. Murphy. Macroeconomic Frameworks: Reconciling Evidence and Model Predictions from Demand Shocks. *NBER Working Paper*, 26365, 2020.

D. Autor, D. Dorn, and G. H. Hanson. On the Persistence of the China Shock. *Brookings papers on economic activity*, Fall:381–447, 2021.

D. H. Autor, D. Dorn, and G. H. Hanson. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American economic review*, 103(6):2121–2168, 2013.

D. Azzopardi, F. Fareed, M. Hermansen, P. Lenain, and D. Sutherland. The Decline in Labour Mobility in the United States: Insights from New Administrative Data, 2020.

R. Barsky, J. Parker, and G. Solon. Measuring the Cyclicality of Real Wages: How Important Is Composition Bias? *Quarterly Journal of Economics*, 109(1):1–25, 1994.

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

G. Basso and G. Peri. Internal Mobility: The Greater Responsiveness of Foreign-Born to Economic Conditions. *Journal of Economic Perspectives*, 34(3):77–98, 2020.

G. Basso, F. D'Amuri, and G. Peri. Immigrants, Labor Market Dynamics and Adjustment to Shocks in the Euro Area. *IMF Economic Review*, 67:528–572, 2019.

S. Basu and C. L. House. Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models. *Handbook of Macroeconomics*, 2:297–354, 2016.

P. Beaudry and J. DiNardo. The Effect of Implicit Contracts on the Movement of Wages Over the Business Cycle: Evidence from Micro Data. *Journal of Political Economy*, 99(4), 1991.

P. Beaudry, D. A. Green, and B. M. Sand. The Declining Fortunes of the Young Since 2000. *American Economic Review*, 104(5):381–386, 2014.

S. Bhattarai, F. Schwartzman, and C. Yang. Local Scars of the U.S. Housing Crisis. *Journal of Monetary Economics*, 119:40–57, 2021.

O. J. Blanchard and L. F. Katz. Regional Evolutions. *Brookings papers on economic activity*, 1992(1): 1–75, 1992.

J. Bound and H. J. Holzer. Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s. *Journal of labor Economics*, 18(1):20–54, 2000.

B. C. Cadena and B. K. Kovak. Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession. *American Economic Journal: Applied Economics*, 8(1):257–90, 2016.

M. Dao, D. Furceri, and P. Loungani. Regional Labor Market Adjustment in the United States: Trend and Cycle. *Review of Economics and Statistics*, 99(2):243–257, 2017.

S. J. Davis and J. Haltiwanger. Labor Market Fluidity and Economic Performance. *NBER Working Paper*, 20479, 2014.

R. A. Decker and J. Haltiwanger. Surging Business Formation in the Pandemic: Causes and Consequences? *Brookings Papers on Economic Activity*, Fall:249–302, 2023.

R. A. Decker, J. Haltiwanger, R. S. Jarmin, and J. Miranda. Declining Business Dynamism: What We Know and the Way Forward. *American Economic Review: Papers and Proceedings*, 106(5): 203–207, 2016.

R. A. Decker, J. Haltiwanger, R. S. Jarmin, and J. Miranda. Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown. *American Economic Review*, 107(5):322–326, 2017.

J. DeWaard, M. Hauer, E. Fussell, K. J. Curtis, S. D. Whitaker, K. McConnell, K. Price, D. Egan-Robertson, M. Soto, and C. Anampa Castro. User Beware: Concerning Findings from the Post 2011-2012 U.S. Internal Revenue Service Migration Data. *Population Research and Policy Review*, 41(2):437–448, 2022.

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

A. Foschi, C. L. House, C. Proebsting, and L. L. Tesar. Should I stay or should I go? The response of labor migration to economic shocks. *Brookings Papers on Economic Activity*, forthcoming.

W. H. Frey. Just Before COVID-19, American Migration hit a 73-Year Low. *Brookings commentary*, 2020.

Y. Gorodnichenko and B. Lee. Forecast Error Variance Decompositions with Local Projections. *Journal of Business Economic Statistics*, 38(4):921–933, 2020.

H. Hyatt, E. McEntarfer, K. Ueda, and A. Zhang. Interstate Migration and Employer-to-Employer Transitions in the U.S.: New Evidence from Administrative Records Data. *Center for Economic Studies*, CES 16-44R, 2018.

Ò. Jordà. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182, 2005.

G. Kaplan and S. Schulhofer-Wohl. Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey. *Demography*, 49:1061–1074, 2012.

M. Kudlyak. How Cyclical is the User Cost of Labor? *Journal of Economic Perspectives*, 38(2): 159–180, 2024.

G. Masnick. Different Data Sources Tell Different Stories About Declining Geographic Mobility. *Harvard Housing Perspective Blog*, 2013.

B. D. Meyer, W. K. C. Mok, and J. X. Sullivan. Household Surveys in Crisis. *Journal of Economic Perspectives*, 29(4):199–226, 2015.

A. Mian and A. Sufi. What Explains the 2007–2009 Drop in Employment? *Econometrica*, 82(6): 2197–2223, 2014.

R. Molloy, C. L. Smith, and A. Wozniak. Internal Migration in the United States. *Journal of Economic Perspectives*, 25(3):173–196, 2011.

R. Molloy, C. L. Smith, and A. Wozniak. Job Changing and the Decline in Long-distance Migration in the United States. *Demography*, 54(2):631–653, 2017.

E. Nakamura and J. Steinsson. Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *The American Economic Review*, 104(3):753–792, 2014.

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

A. Saiz. The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

A Data sources

Here, we present an overview of the data used for our analysis. Details about the data and construction of the Bartik instrument are discussed in a separate section below. For more details on the other instruments used in Table 1, please refer to Foschi et al. (forthcoming).

A.1 Levels of geographic disaggregation

Our analysis examines labor market and population responses at the state level ($N = 48$). 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, and remove observations pertaining to Louisiana in 2005 due to the large displacement of the population following Hurricane Katrina. In Table 1, we also present results for other levels of aggregation, namely commuting zones (CZs, $N = 721$) and counties ($N = 3,049$). For more details on data at these other levels of aggregation, please see Foschi et al. (forthcoming). [19]

A.2 Labor market data

Data on employment, unemployment, and labor force is from the Local Area Unemployment Statistics (LAUS) of the Bureau of Labor Statistics (BLS).We collect data at the county level starting in 1976 and obtain state-level data by summing over the counties in each state.[20] Annual figures are computed as averages of the January-to-December monthly figures. For the extension of our analysis with more historical data (Fig. 8), we replace the employment data from the BLS LAUS with employment series from the State and Area Employment, Hours and Earnings Statistics contained in the BLS's CES, which are available at the state level dating back to 1945. Note that while employment from LAUS refers to the number of people employed by place of residence, employment from the CES measures the number of jobs by place of work; for this reason, we refer to the former as employment and to the latter as jobs .

[19]We rely on the 1990 definition of CZs, and use the crosswalk by Autor et al. (2013) to map counties into CZs. Note that all counties are included in a CZ, and that CZs may cross state lines but never county lines. We also carry out the same analysis for core-based statistical areas (CBSAs): the baseline results are very similar to those for commuting zones and we do not include them, but we use CBSAs for one of the alternative instruments (the defense spending instrument by Auerbach et al. (2020)).

[20]Official data starts in 1990; we obtained unofficial data for 1976-1990 upon request from the BLS.

A.3 Population and migration data

For each county, we take population data from the BEA, which is available at the state level since 1929.[21] Population figures are reported as of July 1: to make them comparable to the annual averages of the labor market data, we use the following adjustment to construct the equivalent of a January-December average: $\overline{Pop}_{i\,t} = 0.25 Pop_{i\,t\quad 1} + 0.5 Pop_{i\,t} + 0.25 Pop_{i\,t\quad 1}$. 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. To further distinguish between population change, domestic and international migration, and births and deaths, we use more detailed Census data from 1990 which contains info on each of these components. For the analyses that require population and net migration by age group, we use the Census population data published by the Survey of Epidemiology and End Results (SEER) of the Cancer Institute, which has a more detailed demographic breakdown with the same time and geographic coverage as the BEA population data.

To study gross migration rates, which require separate info on inflows and outflows for each state, we rely on the Internal Revenue Service (IRS), the American Community Survey (ACS), and the Current Population Survey (CPS). The IRS provides state-level estimates of migration flows since 1976 based on changes in tax return addresses.[22] The ACS has a size of approximately 3 million households in recent years and contains individual-level data; each individual is asked whether they moved their residence from a different state to the current one over the past year. We can accordingly code respondents as movers and non-movers and aggregate them to obtain national gross migration rates. The CPS has a smaller size of 100,000 households, and similarly reports individual movers, from which we obtain aggregate migration figures.

Due to its better quality and long time coverage, we also rely on the IRS data to decompose

[21] This corresponds to the population data produced by the Census Bureau, consisting of decennial counts for Census years, and intercensal estimates for non-Census years. Inputs for the intercensal estimates are vital statistics, data from the Social Security Administration's Numerical Identification File, Medicare enrollment, tax returns from the Internal Revenue Service (IRS), and the American Community Survey (ACS). The IRS tax returns data also form the basis for the IRS migration flow estimates described later.

[22] Migration flows based on the IRS tax returns are one of the key inputs into the production of intercensal estimates. IRS migration data was 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 et al. (2011). Since 2012, the IRS has handled both the production and the publication of the data, and has also introduced some methodological changes. Although these changes can cause concerns about consistency and comparability with earlier numbers (DeWaard et al. 2022), some economists have argued that they have improved the quality of the data, and encourage using IRS data to study trends in gross migration over other other data sources such as the CPS (Hyatt et al. 2018. The similarity of the gross migration trends based on the IRS with those based on the ACS is also reassuring in this respect; the same is true for the similarity between the net migration results in our baseline, which are based on the Census population, and those relying on net migration figures from the IRS, which are almost the same.

gross migration into offsetting migration and absolute net migration. Conversely, thanks to the detail of individual-level characteristics they contain, we use the ACS and CPS to compute migration rates for different demographic groups, based on age, level of education, housing status, and nativity. Finally, we rely on the ACS alone to study net migration by education, housing status, and nativity, which we define as the year-to-year change in population for each state and group.

A.4 Wages, rents, and house prices

To look at the response of wages, rents, and house prices to the Bartik, we put together data from several sources. For wages, we take total wages and salaries from the BEA, and divide them by the number of jobs,[23] also taken from the BEA.

For house prices, we use the house price index (HPI) data from the Federal Housing Finance Agency (FHFA). This is a repeat-sales index based on mortgages purchased or securitized by either Fannie Mae or Freddie Mac. The index is available at the state-level since 1975. To allow comparability across states, we combine the 2000=100 version of the index with median state home values from the 2000 Census to derive state-level house price series.

For rents, we use data on fair-market rents (FMR) from the Department of Housing and Urban Development (HUD). FMRs are estimates of the 40th percentile of gross rents in each county-year, which the HUD uses to determine payment standard amounts for some of its programs. The data is available from 1983. We compute state-level rents as population-weighted averages of the FMRs of the corresponding counties in the state.

B Details on the construction of the Bartik instrument

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

[23] This is an employment series, but follows a number-of-jobs and place-of-work concept like the CES data described in the labor market data section, hence we refer to it as jobs.

that we can rely on approximately 1000 industries overall.

Our sample includes a number of instances when industry classifications changed: there are different versions of the SIC classification, different versions of the NAICS classification, and most importantly, a transition from the SIC to the NAICS system in 1997. We take extra care when constructing the Bartik to ensure that these changes do not cause unexpected behaviors in the instrument. In particular, when constructing the shifts with the QCEW, we rely on SIC data until 1997, the last year when it is 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 et al. (2021) to convert all industry codes to the newest classification whenever we are averaging across years that include a classification change. This ensures that we have always harmonized the classification before computing both averages and changes. We also use these harmonizations to match early-period shares with late-period shifts, and vice-versa, when conducting our variance-decomposition analysis in section 4.1.