Regional labor market adjustment in the United States: trend and cycle

Mai Dao, Davide Furceri, Prakash Loungani*

June 7, 2016

Abstract

We present new evidence on the evolution of labor mobility in the United States over the last four decades. Building on the seminal methodology by Blanchard and Katz (1992), combined with multiple sources of regional population and migration data, we show that interstate mobility in response to relative labor demand conditions is not as high as previously established and has been weakening since the early 1990's. In addition, we find that mobility is counter-cyclical: net migration across regions responds more strongly to spatial disparities in recessions than in normal times. While the declining trend in mobility has been driven by weaker out-migration from states experiencing negative relative shocks, the mobility surge in recessions is mostly accounted for by temporarily stronger in-migration to better-performing states.

JEL Classification Numbers: J61; R23

Keywords: Interstate migration; labor mobility; regional labor markets.

*We would like to thank Olivier Blanchard, Morris Davis, Larry Katz, Justin Wolfers, the editor and three anonymous referees for valuable comments. Helpful feedback from participants at the SED 2015 conference, the HULM 2015 workshop, the EEA 2015 annual meetings, the European Regional Science 2013 conference and various seminars at the IMF is gratefully acknowledged. We thank Monica Cuba and Zidong An for providing excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily represent views of the IMF or IMF policy.
Email: mdao@imf.org (M.Dao), dfurceri@imf.org (D.Furceri), ploungani@imf.org (P.Loungani).

1 Introduction

A high degree of labor mobility has long been considered a distinguishing feature of the U.S. labor market, a view cemented by Blanchard and Katz (1992), henceforth BK. They used a novel method which backed out the response of state-level population to state-specific demand shocks as a residual from the joint adjustment of state employment, unemployment and participation rates. Their method suggested that interstate migration responds quickly and strongly to regional shocks, thus shielding unemployment rates and participation rates from bearing much of the burden of the adjustment.

Building on the BK paper as a starting point, this paper provides a comprehensive analysis of the cyclical and trend behavior of U.S. labor mobility since the mid 1970's. In particular, we are able to take advantage of several data sets that directly measure migration which have become available in the two decades since their paper. Overall, our results paint a picture of U.S. labor mobility that is different in several important ways than the characterization provided by BK and other work in the literature. We also assess how mobility has changed over time and how it behaves during aggregate downturns, including the Great Recession.

Our first key finding is that labor mobility is less important as a cyclical adjustment mechanism in the short-run, relative to changes in unemployment and participation, than suggested in earlier work. We arrive at this conclusion by confronting previous results with interstate net-migration data - available only starting in 1990, right after the end of the BK sample - which provides a direct measure of interstate population movement, as opposed to treating it as a residual. This allows us to test the validity of the BK identification assumption that shocks to regional employment growth reflect relative demand shocks. We find that the BK residual approach and their baseline identification assumption provide estimates for implied labor mobility that are not in accordance with estimates using migration and population data directly. Instead, by incorporating an instrumental variable, the so-called industry mix variable (Bartik, 1991) that is commonly used to measure local demand shocks into the BK framework, we are able to obtain estimates for labor mobility that are statistically equivalent to those obtained from migration data directly. The use of this instrumental variable approach shows that it is primarily the relative unemployment rate, not net migration, that is the main adjustment mechanism in the first two years following a relative shock to state labor demand. When 10 workers in a state experience job loss, while the rest of the country does not, the BK approach implies that 6 of them would migrate out within the first year, leaving only 4 be absorbed by the state-specific unemployment or inactivity pool. Our specification, disciplined by direct migration data, instead implies no statistically significant net out-migration within the first year, with almost all of the shock initially being reflected in the state unemployment rate.

Our second set of findings pertains to a newer literature that documents the long-run decline in gross internal migration rates since the 1980's (see the review in Molloy et al., 2011). There has been no systematic analysis on whether this trend in overall mobility is also associated with smaller (net) migration response to regional disparities over the business cycle, and how this measure of labor mobility behaved during the Great Recession. We fill this gap by establishing the following three results that reveal important patterns in regional adjustment mechanisms.

First, in the last two decades or so, the response of population to regional shocks in the short to medium-run (within the first 5 years) has decreased. Second, the smaller population response to shocks is driven entirely by less population decline in states that experience adverse labor demand shifts, whereas the net population increase in states with favorable labor demand shifts has increased or remained constant (depending on time horizon). Third, despite the trend decline in gross migration rates, the population and migration response to a state-level demand shock increases strongly in recessions, potentially playing a larger role as shock absorber during aggregate downturns than in normal times. Importantly, this counter-cyclical response of population growth is driven primarily by a stronger response of population inflow into states that do relatively better during recessions, while population outflow from states that do relatively worse increases by less and is delayed, occurring toward the end of the recession.

Overall, our results offer a less sanguine view of the ability of U.S. workers to shield themselves from the consequences of adverse shocks than is available in the literature. We show that, particularly in the *short-run*, labor mobility is less important as an adjustment mechanism, and unemployment rates more important, than previously thought to be the case. And while net migration picks up during recessions - despite the trend decline in labor mobility - it benefits regions that do relatively poorly less than others. That said, *long-run* population adjustment still plays an important role in regional responses to shocks, so that the core BK result remains valid.

While there are several papers in the literature that relate to ours, none offers the comprehensive view of U.S. labor mobility that we provide. Beyer and Smets (2015), who use the BK approach to compare US and European labor mobility, also find that extending the BK sample delivers somewhat lower migration response over time. However, they do not use the instrumental variables approach that we take, nor do they use migration data as a cross-check on the results from the BK approach. Other papers study the trend movements in mobility; for instance Partridge and Rickman (2012) also report that the response of net migration to local shocks has declined over time using low-frequency Census data. However, our paper is the first to trace out the response of migration (and other regional labor market variables) to state shocks at business cycle frequency and show the counter-cyclical pattern of these responses. Moreover, we go further than the literature by decomposing these long-run and cyclical patterns to contributions stemming from net in-migration to states with relative positive demand shocks versus out-migration from relatively worse-performing states, hence providing important insights toward understanding the underlying forces.

The rest of the paper is structured as follows. In the next section, we provide some key summary statistics on the persistence and dispersion of regional labor market conditions over time. In section 3, we revisit the panel VAR framework proposed by BK and discuss in detail the identification strategy in section 4. In section 5, we document the cyclical pattern and trend change in regional adjustment over the recent three decades, differentiating between positive and negative state-level shocks, and briefly discuss the underlying mechanisms. Concluding remarks are given in Section 6.

2 Statistical properties of regional employment and unemployment

A prerequisite for labor mobility to absorb and diffuse shocks is the existence of sizeable spatial disparity. An important stylized fact from the BK paper is that U.S. states have been experiencing very different growth rates in employment, and that these different growth rates have been consistently sustained over decades from 1950 to 1990. This section assesses whether this observation still holds in recent years. For this purpose, we split our sample of state-level data and plot average annual employment growth between 1977 and 1994 against the average growth rate between 1995 and 2013 by state, as shown in Appendix Figure A1. The first sub-sample largely overlaps with the second half of BKs sample, during which states showed strong employment growth persistence relative to the preceding decades in the postwar period. By adding two more decades of data, we find that the persistence of state-specific employment growth rates (and similarly, state-specific unemployment rates) still holds - disparities in regional labor markets are therefore long lasting and offer scope for diffusion of shocks through internal migration.

Moreover, we can illustrate the change in spatial disparity by plotting the time-series of the cross-sectional dispersion of state-level employment growth as in Figure 1. We find that dispersion across states has, on average, declined starting in the early 1990's, though it seems to have picked up slightly since the Great Recession. The decline in spatial dispersion has been discussed for example by Kaplan and Schulhofer-Wohl (2013), who argue that it is related to the declining interstate migration rate that occurred during the same time. Interestingly, we consistently observe spikes of high dispersion during periods of recessions. Geographic specialization obviously plays a role for these spikes: as some industries (e.g. construction and auto industries) are more cyclical, that is, sensitive to aggregate shocks than others, a recession hits regions specializing in these cyclical industries (e.g. in Michigan and Nevada) harder, increasing the dispersion of employment across regions. We will explore later in the paper how these spikes in employment dispersion can be derived from increased dispersion of underlying shocks and/or increased employment responses to those shocks.

For the remainder of the analysis, we will look at the joint behavior of state-level labor market variables that cover different labor market statuses. Suppose that each state produces a different bundle of goods, due to different industrial structure, and hence is subject to different shocks or responds differently to aggregate shocks. If a state is hit with a negative relative labor demand shock - that is, relative to the national average - the workers affected either become unemployed, drop out of the labor force, or migrate out of state. We investigate the magnitude and composition of this response by estimating a joint dynamic system in the three state-level variables: employment level, unemployment rate, and labor force participation rate. All labor market outcome variables are taken from various local and national datasets of the Bureau of Labor Statistics (BLS). In particular, state employment and unemployment data are taken from the Local Area Unemployment Statistics (LAUS) dataset from the BLS, which is based among others on CPS and payroll survey data.[1]

For comparability of results, we follow BK in terms of variable definition. The state-relative variables are defined as log deviation from their national aggregates. That is, for employment, e_s is the log employment in state s minus log employment in the US. Consistent with BK, we find that state-relative employment levels are non-stationary as the hypothesis of a unit root cannot be rejected in the majority of the states as well as using panel unit root tests.[2] We therefore use the first difference Δe_s which corresponds to state-relative employment growth. Unlike the relative employment level, the relative log relative employment rate le (equal the negative log relative unemployment rate) and log relative participation rates lp do not exhibit the same persistence and tend to revert to long-term averages.[3]

[1]Appendix Table A1 reports key summary statistics of these state-level data across time and states as well as their detailed sources.

[2]An illustration of this time-series property of relative employment levels is given in Appendix Figure A2, which replicates a similar figure from BK and extends it with two more decades of data. We observe that many states continued their stable regional trend since the post-war period, such as those in the Rust Belt and Mid-Atlantic areas. Others experienced quite large shifts in employment trends, particularly the farm and oil states. Also interestingly, states that were hit hard during the Great Recession in the Sun Belt (Arizona, Nevada, Florida) appear to now be on a permanently lower relative level than their previous trend path, suggesting that the specification of cyclical shocks possibly having permanent effect on relative employment levels (as used by BK and us) is reasonable.

[3]The Im-Pesaran-Shin panel unit root test, allowing for 4 lags, a state-specific constant and a time trend can reject

Overall, we can summarize that the employment growth and unemployment rates across states show strong, albeit weakening, persistence. Moreover, this persistence is related to the persistence of the mean of the employment growth and unemployment rates as opposed to persistent deviations from the means, as the stochastic behavior of both variables displays strong mean-reversion, a feature already documented by BK. Moreover, we document a reduced dispersion of state-level labor market conditions over the last 20 years, stabilizing recently, and with spikes of sharply rising dispersion during each aggregate downturn.

3 The BK estimation approach and results

In this section, we replicate the methodology in BK to estimate the response of state-level labor market variables to a relative shock, adding 23 years of additional data to the original BK exercise to now span 1976-2013. Given the time series properties above, we estimate a system of panel VAR equations as follows:

$$\begin{aligned}
\Delta e_{st} &= \alpha_{s10} + \alpha_{11}(L)\Delta e_{s,t-1} + \alpha_{12}(L)le_{s,t-1} + \alpha_{13}(L)lp_{s,t-1} + \epsilon_{set}, \\
le_{st} &= \alpha_{s20} + \alpha_{21}(L)\Delta e_{s,t} + \alpha_{22}(L)le_{s,t-1} + \alpha_{23}(L)lp_{s,t-1} + \epsilon_{sut}, \\
lp_{st} &= \alpha_{s30} + \alpha_{31}(L)\Delta e_{s,t} + \alpha_{32}(L)le_{s,t-1} + \alpha_{33}(L)lp_{s,t-1} + \epsilon_{spt}. \qquad (1)
\end{aligned}$$

We pool all states while allowing for state-specific constants, thus estimating the dynamics of the average state. We include two lags for each variable, following BK, and to keep sufficient degrees of freedom for estimation with shorter sub-samples, though extending up to four lags does not change the estimates substantially. This identification strategy assumes that current unexpected changes to state-relative employment growth within the year primarily reflect movements in regional labor demand. This assumption allows us to estimate the dynamic effects of a 1 percent shock to labor demand in a typical state on its relative unemployment rate, labor participation rate,

the hypothesis of a unit root for the relative log employment rate (the negative of the relative log unemployment rate) le and relative log participation rate lp at the 5 and 10 percent level of significance respectively.

and as a residual, the net population change of the state. This is because in any period, we can decompose the change in the relative log employment level de (where d denotes the change relative to pre-shock baseline) into:

$$de = dle + dlp + m, \tag{2}$$

where m stands for the implied log change in state-level civilian, non-institutional working-age population head count, which can be driven by mortality, incarceration, immigration from abroad, and most importantly for our exercise, net migration across state.

There are several ways to estimate the system of equatons (1). Given the identification assumption that current shocks to employment growth are driven by labor demand only, $\Delta e_{s,t}$ is weakly exogenous in the equations for le and lp and the system can be consistently estimated by OLS equation-by-equation, which is the estimation we use. This is identical to transforming the system to a reduced form VAR and ordering employment growth first. We also use panel GMM to estimate the system to control for potential inconsistency of OLS caused by fixed effects in the presence of lagged dependent variable. Given the long time series, the difference in estimation results is marginal (results available upon request).

The results imply that a negative 1 percent shock to labor demand in a state raises its unemployment rate by 0.2 percentage points and lowers the participation rate by 0.3 percentage points relative to the national average in the first year, with the effect peaking at 0.3 and -0.4 percentage points after 2 years respectively (and symmetrically for positive shocks).[4] The effect on the relative employment level peaks after four years at -1.7 percent, before decreasing gradually to a long-run value of around -1.2 percent. The response of relative population growth is derived as a residual and amounts to a net population decline of 0.4 percent (of initial working-age population) in the first year and 0.6 percent in the second. Over the long run, employment *growth*, as well as unemployment and participation rates revert to the pre-shock level, while the employment level is permanently changed. That is, interstate population adjustment following the temporary regional

[4] Appendix Figure A4 shows the complete set of impulse responses estimated using OLS to a negative 1 percent shock to relative labor demand (and symmetrically to a positive 1 percent shock).

shock drives permanent changes in relative employment levels. It is also instructive to translate the changes from rates to number of workers. Of every 10 workers that lose employment, 2 workers become unemployed, 2 drop out of the labor force, and 6 workers migrate out of state within the first year following the shock. Compared to the original BK results, differences purely due to data updates are not large: subject to the same negative shock, participation falls more (0.4 instead of 0.3 percentage points at trough) and therefore, migration responds somewhat less in the short-run (net out-migration of 0.8 instead of 1 percent of population by year 3).

Based on these findings, a well-known conclusion from the BK paper is that most of the (short- and long-run) response to regional shocks occurs through net migration. Furthermore, the apparent stability of the BK results over time suggests that this pattern of adjustment remained roughly unchanged in the last 20 years. However, a crucial assumption underlying the BK results for this conclusion to hold is that shocks to employment growth across states are entirely driven by variation in state-specific labor demand. We devote the following section to examining the validity of this assumption.

4 Endogeneity of state labor demand shocks

4.1 Test of OLS identification assumption

In this section, we take a step back to test the identification assumption of BK that was used for the OLS estimation above, as well as by many other ensuing studies of labor mobility (see e.g. Decressin and Fatas, 1995; Jimeno and Bentolila, 1998). The crucial assumption is that unexpected shocks to relative employment growth, that is ϵ_{set} in the first equation of the system (1), are purely state-relative labor *demand* shocks. To test this assumption, we use as instrumental variable (IV) the so-called industry shift or industry mix variable, first proposed by Bartik (1991) and subsequently used extensively in the urban/regional economics literature. This variable measures the predicted employment growth in each state based on the state's industrial composition of employment and the overall employment growth of each industry. More precisely, the industry mix variable $imix_{st}$ is defined as:

$$imix_{s,t} = \sum_{j=1}^{J} \left[\bar{\theta}_{sjt} \Delta \ln(\bar{e}_{jt,-s}) \right], \tag{3}$$

where the state-specific industry share of employment $\bar{\theta}_{sjt}$ is taken as a 5-year moving average to avoid endogeneity with respect to current regional labor market conditions, and aggregate industry employment growth $\Delta \ln(\bar{e}_{jt,-s})$ is the growth rate of each industry j in all US states excluding state s. The state-level industry employment shares as well as the industry-level employment growth rates ($\Delta \ln(\bar{e}_{jt})$) are taken from the Bureau of Economic Analysis (BEA) Regional Economic Accounts. The industries j are based on 20 2-digit code SIC industries up until 2000, and 20 2-digit code NAICS industries starting in 2001, both covering full and part-time jobs in the entire non-farm private sector. The identification relies on the pre-determined production structure of each state and each industry's national out-of-state growth rate, which are both arguably uncorrelated with state-specific labor supply shocks.[5]

Because we are interested in states' relative labor market outcomes (relative unemployment, relative participation), and because net-migration responds to relative, not absolute labor market conditions (see BK and other models of spatial equilibrium such as Roback, 1982), we take deviations of $imix$ from their national averages in each year to obtain measures of relative labor demand changes ($rimix_{s,t}$):

$$rimix_{s,t} = imix_{s,t} - \overline{imix_t}.$$

Using $rimix_{s,t}$ to instrument for $\Delta \ln e_{s,t}$ in the equations for relative employment rate le and participation rate lp from the system of equations (1), we obtain the 2SLS results summarized in column 3-4 of Table 1, with OLS results presented for comparison in column 1-2.

The first stage regression results show, as also illustrated in Appendix Figure A5 for different sub-samples, that the prediction power of the industry mix instrument for state-level employment

[5] Appendix Table A2 provides a snapshot of the distribution (in 2012) of employment across the different industries, as well as the variation in each industry's employment share across states.

growth is strong (reflected in the large, positive 1st stage coefficient and high F statistics). Using this IV, the second stage result in column 3 reveals a much stronger response of the state-relative employment (or unemployment) rate to state-specific labor demand shocks than do OLS results: a 1 percent negative labor demand shock reduces the employment rate by 0.8 instead of 0.2 percent as with OLS. The Hausman test therefore clearly suggests a rejection of the exogeneity assumption in the OLS regression used by BK.[6]

Results for the participation rate equation using the industry mix variable also lead to rejection of the OLS identification assumption (column 4). The response of state-relative participation rate is in fact smaller using IV than OLS: a 1 percent negative employment shock reduces the participation rate by 0.1 percent instead of 0.4 percent with OLS and is not statistically significant. [7]

Why OLS and 2SLS identification assumptions would yield different results? We believe that the assumption that innovations in relative employment growth reflect purely changes in relative labor demand is likely to be violated. In other words, shocks to relative labor supply may also affect relative employment growth in the same year. For example, migration shocks (triggered by events abroad) may affect labor supply in some states (such as border states), leading to deviation of those states' unemployment *and* employment growth rates from the national average within the same year of the shock. Also, the relative labor force participation rate across states can change if states differ in their age composition and there are abrupt shifts in the size of cohorts entering working-age or retirement age (as has been the case with the retirement of baby-boomer cohorts in recent years). Estimating a structural model of regional labor markets, Partridge and Rickman (2003) find that such relative supply shocks can account for a substantial share of variation in state employment growth from year to year. In fact, the sign of the OLS bias we find in Table 1 is consistent with

[6] In their paper, BK also carry out a similar 2SLS regression and conclude that their OLS identification is robust. We replicate their result and conclude that the reason for this discrepancy with our finding is the short sample of 10 years (1978-1988) for which the industry-mix IV was available at the time. This severely limits the degrees of freedom, exacerbates the fixed-effect induced bias in the panel regression, and makes the IV only a weak instrument for contemporaneous employment growth in the first stage, while biasing estimates towards OLS in the second stage.

[7] We confirm the results using an alternative IV that picks up exogenous changes to state-level labor demand in oil and gas extraction industries triggered by changes to the aggregate oil price, an identification strategy that has also been used in e.g. Saks and Wozniak (2011) and Gallin (2004). The results confirm the findings using the industry-mix IV: OLS underestimates the response of state-level unemployment rates and over-estimates the response of participation rates to state-level shocks (see Appendix B).

relative labor demand and supply shocks being confounded. If part of the innovation to relative employment growth reflects shocks to relative labor supply, then we should expect a negative correlation between contemporaneous employment growth and the residual in the employment rate (le) equation, as stronger labor supply temporarily increases the unemployment rate and reduces the employment rate. This source of endogeneity would bias the OLS estimate in the le equation toward zero, which is exactly what we find. At the same time, positive shocks to labor supply would increase the participation rate, resulting in a positive correlation of contemporaneous employment growth and the residual in the lp equation. The same source of endogeneity would therefore cause OLS estimate to be biased upward in the lp equation, which our estimates in Table 1 also confirm.

To sum up, we find that the BK identification assumption for relative labor demand shocks is not supported by the data. By confounding relative labor demand and supply shocks, the BK methodology underestimates the response of the relative unemployment rate and overestimates the response of the participation and population in the short-run.[8]

4.2 A new framework to estimate regional adjustment

In light of the preceding results on the endogeneity of contemporaneous employment growth, we provide in the following a modified version of the BK framework that represents a reduced form of the 2SLS estimation of the previous section. To trace the joint dynamic response of each labor market variable to a regional labor demand shock using the industry mix variable, we estimate the

[8]There could also be an OLS attenuation bias due to measurement error in state-level employment growth (especially for small states) that the IV can address, as it only picks up the variance in the "signal" component of the potentially mis-measured employment growth variable. Remaining measurement error of the "signal" in the IV is also likely to be smaller, as it is constructed using averaged data across states and years, hence largely smoothing out i.i.d. measurement errors.

following reduced-form VAR system with $rimix$ being an exogenous forcing variable:[9]

$$\begin{aligned}
\Delta e_{st} &= \alpha_{s10} + \alpha_{11}(L)rimix_{s,t} + \beta_{11}(L)\Delta e_{s,t-1} + \beta_{12}(L)le_{s,t-1} + \beta_{13}(L)lp_{s,t-1} + \epsilon_{set}, \\
le_{st} &= \beta_{s20} + \alpha_{21}(L)rimix_{s,t} + \beta_{21}(L)\Delta e_{s,t-1} + \beta_{22}(L)le_{s,t-1} + \beta_{23}(L)lp_{s,t-1} + \epsilon_{sut}, \\
lp_{st} &= \beta_{s30} + \alpha_{31}(L)rimix_{s,t} + \beta_{31}(L)\Delta e_{s,t-1} + \beta_{32}(L)le_{s,t-1} + \beta_{33}(L)lp_{s,t-1} + \epsilon_{spt}. \qquad (4)
\end{aligned}$$

To illustrate the difference between these reduced-form IV estimates, henceforth RFIV-VAR, and those pertaining to the OLS-VAR from the original BK specification (1), Figure 2 plots the response of each labor labor market variable, including the net change in population, to a 1 percent negative labor demand shock resulting from both models (responses are symmetrical for a positive shock). That is, we compare the dynamic response to a 1 percent negative shock to $\Delta e_{s,t}$ obtained from the OLS-VAR model (1) (setting $\epsilon_{set} = -0.01$), versus the response to a shock of same magnitude to $\Delta e_{s,t}$ from the RFIV-VAR model (4) (setting $\Delta rimix_{s,t} = -0.01 * (\alpha_{11}(0))^{-1}$).

Consistent with 2SLS estimates in Table 1, the response of the state-relative unemployment rate to a given labor demand shock is much stronger using RFIV-VAR than OLS-VAR in the first two years following the shock, whereas the participation rate responds less at all horizons. The implied net population response is therefore weaker under RFIV-VAR than under OLS-VAR: a 1 % labor demand shock leaves the working-age population unchanged instead of reducing it by 0.4 % through net migration within the first year. Similarly, the long-term adjustment through net population change is weaker by a third with RFIV-VAR than OLS-VAR, leading to a smaller total employment change (around 0.8 % instead of 1.2 %). Translated to absolute changes, when a relative negative shock causes 10 workers to lose employment, the OLS-VAR estimates imply that 6 of them will migrate out of state, 2 become unemployment and 2 drop out of the labor force in the first year. Instead, the RFIV-VAR estimates imply no population change in the first year, but

[9] We use the reduced form instead of 2SLS regression in the system for two reasons. Conceptually, this allows us to either condition results to a unit change in $rimix$, corresponding to a given shift in ex-ante relative labor demand, which will prove useful for analyzing changing sensitivity over time, or to a unit change in ex-post employment growth, which will prove useful for comparison with OLS. Econometrically, reduced form estimation avoids the small-sample bias of 2SLS as formalized in Chernozhukov and Hansen (2008).

9 workers becoming unemployed and 1 dropping out of the labor force. After ten years, the same shock would have led 20 people to move out under OLS-VAR compared to 15 under the RFIV-VAR.[10] Overall, population adjustment across states does not act as quickly and strongly to smooth out spatial disparities as previously thought, allowing the pass through of ex-ante disparities to regional unemployment to be on average three times higher than estimated by BK in the short-run.[11]

4.3 Validation of results with net migration data

So far, the implied population response was backed out from the response of the employment and participation rates (as they jointly pin down the change in working-age population). We expect the population response to be primarily driven by net-migration response across states, as the differential in adult mortality, incarceration and foreign immigration are less likely to respond immediately to state-level demand shocks. This approach is particularly useful as sufficiently reliable migration data is not available for long time periods. However, several datasets containing information on geographic mobility became available after the original BK paper. It is therefore interesting to compare the derived response with one that is estimated using migration data directly.

The main migration dataset we use is the annual *State Population Estimates and Demographic Components of Change* data from the US Census Bureau's Population Estimates Program (Census PEP). The annual population estimates start with the decennial census data as benchmark and add annual population component of change data, that is births, deaths, internal migration, immigration, emigration, and Federal (armed forces and civilian) movements, which derive from various governmental administrative records and census distributions. In particular, state-level net domestic migration, our variable of interest, is derived by computing the net migration rate implied by

[10] Note that although results in Figure 2 for cumulative net migration after 2-5 years appear similar across the two specifications, this is in fact not because OLS and IV estimates of the system converge in the medium term. Instead, while OLS underestimates the unemployment response, it overestimates the participation response relative to the reduced form, so that the net migration response, derived as residual from $de - dle - dlp$, happens to be similar in years 2-5 (with employment response de normalized to be equal in year 1).

[11] That is, unemployment increases by 5 times more in the first, 3.4 times more in the second, and 1.1 more in the third year, thus on average 3.2 times more in the first 3 years using the RFIV-VAR relative to OLS-VAR estimates.

the share of tax filers and dependents (i.e. exemptions) who changed addresses between any two tax filings based on IRS supplied Federal tax returns for the population 64 years and younger, and from Medicare enrollment data for the population 65 years and older. This methodology to account for domestic migration (and separately, for international migration) was only introduced for the post-1990 population estimates, with the previous years' estimates only accounting for births and deaths and other components of change lumped into one residual. The available sample of state-level domestic net migration data therefore starts in 1991. This measure of state-level net population change excludes variations across states due to mortality, incarceration and international immigration, and hence is closely related to the labor mobility concept we are interested in.

To validate the OLS-VAR identification from model (1) using migration data, we estimate the following equation with the state-level net migration rate as the dependent variable, and with relative labor demand identified by unexpected relative employment growth Δe_{st}, same as in the OLS-VAR identification above:

$$m_{st} = \alpha_s + \gamma_t + \beta(L)m_{s,t-1} + \gamma(L)\Delta e_{s,t} + \delta_1(L)le_{s,t-1} + \delta_2(L)lp_{s,t-1} + \epsilon_{s,t}, \tag{5}$$

where m_{st} is the state-level net migration rate, i.e. annual domestic net migration flow as a share of state population at the beginning of the year, in deviation from a state-specific linear trend to account for long-run trends in state-specific migration evolution (due to e.g. amenities, industry agglomeration) as well as aggregate mobility trends (in particular the secular overall decline in migration documented in the literature). We also allow for state-specific intercepts which capture the effect of time-invariant factors, as well as time fixed effects to control for cyclicality in residual migration (see Saks and Wozniak, 2011 and discussion below). We include the other lagged explanatory variables from the OLS-VAR system (1), so that a contemporaneous change to Δe_{st} in equation (5) is the same unexpected innovation as the one captured in the OLS-VAR. We compute the cumulative response of net migration to a given shock to Δe_{st} estimated directly by equation (5) and compare it to the response backed out from the OLS-VAR system (1). The paths of Δe_{st},

le_{st} and lp_{st} used for computing the response of m_{st} from equation (5) are calibrated to exactly match the respective path from the OLS-VAR system (1).

To perform the same cross-validation exercise for the RFIV-VAR model (4), we estimate the following equation using the interstate net migration data as dependent variable:

$$m_{st} = \alpha_s + \delta_t + \beta(L)m_{s,t-1} + \gamma(L)rimix_{s,t} + \epsilon_{st}, \tag{6}$$

where the relative labor demand shock is identified by relative employment growth predicted by a state's relative industry mix predicted employment growth $(rimix)$. Furthermore, two lags of the dependent variable and the exogenous variable are allowed to be consistent with the RFIV-VAR specification. We simulate the cumulative response of net migration implied by the estimated equation (6) and compare it with the cumulative response of net population change backed out from the RFIV-VAR system (4) following a shock to $rimix$ of the same size.

Note that for these cross-validation exercises, we re-estimate each VAR system (1) and (4) using the same sample period as is used for the migration equations (5) and (6), namely 1991-2013. Panel A and B in Figure 3 present the cross-validation results for the OLS and IV specification respectively. We can see in panel A that the OLS identification yields a large discrepancy between the data and VAR-implied responses of state population to the same labor demand shock. The discrepancy widens with the time horizon but is large both in the short and long term. In contrast, panel B shows that the identification of state-relative labor demand shocks using $rimix$ leads to a very close result between the net population response derived from the VAR model and that estimated with Census/IRS migration data directly, particularly in the short and medium term.

Though Figure 3 clearly illustrates the advantage of our new estimation framework compared to the BK approach in terms of external consistency, we would also like to formally test for the degree of this external consistency. In the following, we develop a test for the over-identification of net migration response implied by data and the VAR systems. For the OLS identification, we stack the OLS-VAR system of equations (1) on the single equation for net migration (5) and estimate an augmented system jointly:

$$
\begin{aligned}
\Delta e_{s,t} &= \alpha_{s10} + \alpha_{11}(L)\Delta e_{s,t-1} + \alpha_{12}(L) le_{s,t-1} + \alpha_{13}(L) lp_{s,t-1} + \epsilon_{set}, \\
le_{s,t} &= \alpha_{s20} + \alpha_{21}(L)\Delta e_{s,t} + \alpha_{22}(L) le_{s,t-1} + \alpha_{23}(L) lp_{s,t-1} + \epsilon_{sut}, \\
lp_{s,t} &= \alpha_{s30} + \alpha_{31}(L)\Delta e_{s,t} + \alpha_{32}(L) le_{s,t-1} + \alpha_{33}(L) lp_{s,t-1} + \epsilon_{spt}, \\
m_{s,t} &= \alpha_{s40} + \delta_t + \beta(L) m_{s,t-1} + \alpha_{41}(L)\Delta e_{s,t} + \alpha_{42}(L) le_{s,t-1} + \alpha_{43}(L) lp_{s,t-1} + \epsilon_{smt}. \quad (7)
\end{aligned}
$$

The system estimates allow us to test for equality of implied and direct migration response by testing for restrictions on parameter values across equations. For example, in the first year following a shock to relative labor demand Δe_{st}, this restriction amounts to the null hypothesis:

$$
H_0 : \quad 1 - \alpha_{21}^0 - \alpha_{31}^0 = \alpha_{41}^0,
$$

where the superscripts 0 index the first coefficient of each lag polynomial. Beyond the first year, the test statistics quickly become highly non-linear, so we restrict the test to the first 3 years after the shock and resort to the delta method.

Following the same principle, we test for over-identification of net migration response within the IV framework by stacking the RFIV-VAR system (4) on the single equation (6) and jointly estimating the augmented system:

$$
\begin{aligned}
\Delta e_{s,t} &= \alpha_{s10} + \alpha_{11}(L) rimix_{s,t} + \beta_{11}(L)\Delta e_{s,t-1} + \beta_{12}(L) le_{s,t-1} + \beta_{13}(L) lp_{s,t-1} + \epsilon_{set}, \\
le_{s,t} &= \beta_{s20} + \alpha_{21}(L) rimix_{s,t} + \beta_{21}(L)\Delta e_{s,t-1} + \beta_{22}(L) le_{s,t-1} + \beta_{23}(L) lp_{s,t-1} + \epsilon_{sut}, \\
lp_{s,t} &= \beta_{s30} + \alpha_{31}(L) rimix_{s,t} + \beta_{31}(L)\Delta e_{s,t-1} + \beta_{32}(L) le_{s,t-1} + \beta_{33}(L) lp_{s,t-1} + \epsilon_{spt}, \\
m_{s,t} &= \alpha_s + \gamma_t + \beta(L) m_{s,t-1} + \gamma(L) rimix_{s,t} + \epsilon_{st}. \qquad (8)
\end{aligned}
$$

The resulting cross-equation restriction for equality of net migration response in the first year is given by:

$$H_0: \quad \alpha_{11}^0 - \alpha_{21}^0 - \alpha_{31}^0 = \gamma^0,$$

We use the Census/IRS migration data as the direct measure for $m_{s,t}$ in the stacked systems above. The resulting Chi-squared test statistics and the p-value under the null hypothesis for the first 3 years after a given shock to relative labor demand are summarized in Table 2. The test results confirm the visual conclusion from Figure 3. While the OLS identification can be rejected at confidence levels of 99 percent or higher at all three time horizons, the IV identification yields estimates for implied migration responses that are statistically indistinguishable from directly estimated ones.

In addition to the Census/IRS net migration data, we also use state-level working-age population growth data from LAUS-BLS as well as working-age migration data from the American Community Survey (ACS) to externally validate the residual migration estimates from the VAR models. Further discussion of these alternative data sources, including their comparability with the Census/IRS migration data, as well as results of these additional validation exercises are summarized in Appendix C. All datasets and tests unanimously support our identification strategy adopted in the RFIV-VAR system (4) and strongly reject the original BK identification assumption. These new estimates have important implications for the dynamics of regional adjustment. Contrary to the long-established results in BK, it is primarily the relative unemployment rate, not net migration, that absorbs affected workers in the first two years following a negative shock to state labor demand. Migration acts as a much weaker mechanism for spatial diffusion of shocks in the short-run and leads to less agglomeration effects in the long-run than previously thought.

5 The evolution of regional adjustment

One important purpose of the paper is to document whether patterns and channels of regional adjustments change over time. The migration literature has long documented a decline in interstate migration rates starting in the 1980's, but does this decline also imply a reduced sensitivity of migration to spatially disparate shocks? Figure 4 plots the implied migration response to a 1 percent shock to predicted employment growth as derived earlier, separately for three different samples: the BK sample of 1976-1990, the subsequent sample up to the crisis 1991-2007, and finally 1991-2013 which includes the crisis years. We overlap the last two sub-samples to have sufficiently long time-series necessary for reliable VAR estimates. This presentation of the data suggests that migration sensitivity to regional shocks, both in the short and long-run, has been strongly decreasing since the 1990s, yet seems to have risen during the Great Recession and its aftermath.

In short, there is suggestive evidence of changes in migration responsiveness over time as well as a shift during the latest recession. We study this more rigorously by developing an estimation strategy to track these changes from year to year, while also distinguishing between relative positive versus relative negative state-level shocks. Indeed, symmetry in regional adjustment may not hold, as documented for regional adjustment to long-term changes by Notowidigdo (2014). If this is also the case for adjustment to cyclical disparities, then differentiating relative positive from negative shocks can offer important insights toward understanding the patterns we have documented. Does the gradual weakening of interstate population response to a given relative shock result from weaker net migration to relatively better performing states, or is it driven by less net migration from worse performing ones over time? Similarly, the counter-cyclical pattern of migration response could be driven by more people leaving states with worse prospects, or more people moving to states with better prospects during recessions, or both. In the following, we "unpack" the main results along these dimensions.

5.1 Positive vs. negative relative shocks

In the following, we modify the main estimation framework in the equation system (4) by allowing for differential response of the system to relative positive vs. relative negative state labor demand shocks:

$$\begin{aligned}
\Delta e_{st} &= \alpha_{s10} + \alpha_{11}^{+}(L) rimix_{s,t}^{+} + \alpha_{11}^{-}(L) rimix_{s,t}^{-} + \beta_{11}(L)\Delta e_{s,t-1} + \beta_{12}(L) le_{s,t-1} + \beta_{13}(L) lp_{s,t-1} + \epsilon_{set}, \\
le_{st} &= \beta_{s20} + \alpha_{21}^{+}(L) rimix_{s,t}^{+} + \alpha_{21}^{-}(L) rimix_{s,t}^{-} + \beta_{21}(L)\Delta e_{s,t-1} + \beta_{22}(L) le_{s,t-1} + \beta_{23}(L) lp_{s,t-1} + \epsilon_{sut}, \\
lp_{st} &= \beta_{s30} + \alpha_{31}^{+}(L) rimix_{s,t}^{+} + \alpha_{31}^{-}(L) rimix_{s,t}^{-} + \beta_{31}(L)\Delta e_{s,t-1} + \beta_{32}(L) le_{s,t-1} + \beta_{33}(L) lp_{s,t-1} + \epsilon_{spt},
\end{aligned} \tag{9}$$

where

$$\begin{aligned}
rimix_{s,t}^{+} &= rimix_{s,t} \text{ if } rimix_{s,t} > 0, \ 0 \text{ otherwise} \\
rimix_{s,t}^{-} &= rimix_{s,t} \text{ if } rimix_{s,t} < 0, \ 0 \text{ otherwise}
\end{aligned}$$

Figure 5 plots the response of state-level population to positive versus negative labor demand shocks, as backed out from the asymmetric model in equation system (9). There is clear evidence for a strong asymmetric response at all time horizons. For example, one year after the shock, population adjusts by 0.6 percent if the relative labor demand shock is positive as compared to 0.2 percent if it is negative, with the difference widening in subsequent years and being strongly statistically significant ($p < 0.01$ by year 3). That is, a state that experiences a 1 percentage point higher employment growth rate relative to the national average (and its own historical average), as predicted by its industrial specialization and aggregate industry demand, attracts a net population inflow that is three times stronger than the net population outflow from a state experiencing a negative shock of equal magnitude. This asymmetry result is consistent with Notowidigdo (2014), who finds that positive shifts in labor demand across MSA's that persist over decades trigger stronger population gains than negative labor demand shifts reduce population. We show that the asymmetry also holds for cyclical shifts in relative labor demand and annual population adjustment across states. The result could also reflect a lack of so-called migration directedness documented by Yagan (2014) using individual-level data, that is, people who move into the better-performing states do not disproportionately come from the worse-performing states.

5.2 Decomposition of the adjustment pattern with expanding window regression

Having shown that positive changes to relative labor demand have a stronger effect on population adjustment than negative ones, we are now ready to trace out the evolution of regional adjustment to positive and negative shocks over time. To this end, we carry out a sequence of expanding window regressions of the assymetric RFIV-VAR equation system (9) starting with the base sample 1976-1990 (the BK sample). We then expand the sample by adding one year at a time and re-estimate the RFIV-VAR. The difference in estimates between any consecutive expanding windows reflects how the last year of observation changes the estimated average dynamics. This allows us to construct annual changes between 1990 and 2013 to any statistics of interest.[12] After estimating a VAR system for each sub-sample, we calculate the response of net population change (proxy for net migration) to a 1 percentage point change in relative (positive and negative) predicted employment growth ($\Delta rimix = +/-0.01$), at different time horizons. To enhance representativeness and keep the estimation from being overly influenced by small states with big shocks in the marginal year, we weight the observations by state-level population (averaged over the sample period). The resulting sequence of estimated migration response to a relative positive versus negative labor demand shock of equal magnitude, both in the short and long run, are presented in Figure 6, with the overall response also plotted for comparison.[13]

Two main results stand out. First, there is an overall downward trend in migration response, especially in the short- and medium-run (see also Appendix Figure D12), which is driven overwhelmingly by a declining migration response to relative negative changes to labor demand. The estimated response to a negative relative shock estimated over the whole sample until 2013 amounts to less than a third the response estimated until 1990. By contrast, migration response into states with relative positive labor demand shifts has been either stable (over the short run) or even increasing over time when measured over the long run. Hence, states that perform better attract

[12] These methods have been widely used in the finance literature, in particular for forecasting purposes. See e.g. Pesaran and Timmerman (2002).

[13] The un-weighted series delivers largely the same result, but is somewhat more volatile.

more population inflow than states that perform worse lose to outflows; moreover, this asymmetry has been widening over time since the early 1990's. Secondly, the increased migration responsiveness during the Great Recession is driven primarily by increased responsiveness of migration into good states, while the outflow intensity from bad states only picked up toward the end of the recession and early recovery.[14]

In Appendix D, we show that unlike for population/migration, a counter-cyclical sensitivity is not consistently observed for the other margins of labor market adjustment (employment and participation rates). At the same time, adjustment of all three margins to a given ex-ante relative shock has been weakening since the early 1990s, implying smaller ex-post variation in employment growth across states subject to similar variation in ex-ante relative demand shocks.[15] Both results are consistent with the earlier observation (from Figure 1) that variation in employment growth has been declining on average until recently, but spikes up during aggregate downturns. Moreover, the declining trend does not appear to be reflected to the same extent in a gradually declining variation of underlying ex-ante shocks to state-relative demand: The cross-sectional dispersion of the exogenous underlying shock $rimix$ - while on average higher in recessions - does not exhibit a declining trend since the early 1990's (see Appendix Figure A8). Less responsiveness to similarly dispersed shocks thus appears to drive the declining cross-section dispersion of employment growth over time, while stronger migration response to more dispersed shocks drive the higher dispersion of employment growth in recessions.[16]

Finally, note that our new result on counter-cyclical migration response is consistent with the previous finding in the literature that gross migration for reasons other than spatial labor market arbitrage is pro-cyclical. Saks and Wozniak (2011) find that *after* controlling for relative labor market conditions between any pair of states, the residual component of state-level gross migration is

[14]There is also a small compositional effect underlying the counter-cyclical pattern, see Appendix C.

[15]The declining trend in participation adjustment since the early 1990s is consistent with a declining overall aggregate labor force participation rate in the US due to aging demographics and hence less mobility into and out of the labor force as older workers' participation rate is less cyclical (see Balakrishnan et al., 2014).

[16]We have also extensively studied the response of wages and its evolultion but do not find statistically significant results for the years after the BK sample, likely due to measurement and compositional problems of underlying wage data (see online Appendix H).

pro-cyclical, rising in expansions and declining in recessions. Interestingly, our results suggest that this is also the case for residual *net* migration. As the effect of the relative labor demand variable $rimix$ is substantially higher in recessions, and because dispersion of $rimix$ is somewhat higher in recessions than expansions (see Appendix Figure A8), our results in fact imply that a higher share of cross-sectional variance of net migration rate is explained by relative labor market conditions in recession. As a consequence, determinants of net migration other than relative labor market conditions (such as amenities, life-cycle) play a smaller role during recessions than expansions, consistent with the Saks and Wozniak (2011) finding for gross migration.[17]

5.3 Cross-validation of the counter-cyclical migration response

Next, to further validate our findings from the expanding window regressions, we assess whether the cyclical pattern of mobility is also reflected in direct measures of interstate migration. As migration data typically have too short a time series to conduct an expanding window regression, we test the counter-cyclical response using those by estimating the following equation with business cycle interaction terms:

$$m_{s,t} = \alpha_s + \gamma_t + \rho m_{s,t-1} + \gamma_1 D(Exp)_t \times rimix_{s,t-1} + \gamma_2 D(Rec)_t \times rimix_{s,t-1} + \epsilon_{s,t}, \quad (10)$$

where $D(Rec)_t$ stands for a dummy variable that equals one if year t contains one or more quarters of NBER-dated recessions, and $D(Exp)_t$ is a dummy variable for years without any recessionary episodes. Using the parameter estimates for γ_1 and γ_2, we can thus test whether the response of net migration is different during aggregate recessions versus normal times. We do not estimate the dynamic path for the response as done above due to the short time series of the data, and instead focus on the response to one-year lagged relative shocks, given that previous results strongly indicate that most of the adjustment materializes 1-2 years after the shock.[18]

[17] Appendix F derives a decomposition of cross-sectional variance of net migration over time, as illustrated in Appendix Figure A9, and confirms that residual factors orthogonal to relative labor demand explain less of the variation of net migration across states during the Great Recession than other years.

[18] We confirm that the contemporaneous response is close to zero.

Using Census interstate migration data, column (1) of Table 3 confirms again that population responds with at least 1 year lag after a relative shock to labor demand across states occurs. Parameters for the business cycle interaction terms according to equation 10 are estimated in column (2) to be positive and statistically significant. Quantitatively, it implies that the population response is more than three times as large during a recession. While a change in relative labor demand of 1 percent increases population through net migration by roughly 0.3 percent after 1 year and 0.5 percent in the long run during normal times, the response during recessions is 1 percent after 1 year and 1.7 percent in the long-run. As 22 years of data may be too short to alleviate the inconsistency of estimation with fixed effects in dynamic panels, we also present results without the lagged dependent variable in column (3), which can only estimate the short-run effect. The estimates for the response after 1 year are virtually unchanged, 0.3 percent in normal times versus 1.2 percent in recessions, with the differential between recession and expansions being statistically significant at less than 1 percent.

So far, we have only used state-level data, as opposed to more granular county or metropolitan area data, to estimate regional adjustment dynamics. The main reason is that for more granular geographic areas, some of the labor market variables are either not available or not for sufficiently long periods of time to conduct the VAR-type estimation.[19] Therefore, we have so far carried out all baseline estimations using state-level labor force statistics only and used interstate migration data to cross-validate the results. However, without tracing out the complete dynamic response with the VAR, it is possible to use substate net migration data directly to test whether short-run population change across these areas responds counter-cyclically to local shocks, providing yet another cross-validation of our key results. For this purpose, we use the Census Bureau PEP

[19] For example, the BLS does publish substate data as part of the LAUS program, but for Metropolitan Statistical Areas, only limited time series can be consistently used as delineation of MSA borders are changed periodically (last two times in February 2013 and December 2009). Apart from the short sample issue, the LAUS data does not provide statistics on labor force participation rates for substate geographic areas, preventing us from using such substate data for the VAR analysis. Data aggregated from the CPS into *Geographic Profile of Employment and Unemployment* from the BLS does have statistics on MSA-level participation rates (going back to 2005), but in turn does not have data on employment level. Estimates of levels are not provided for substate areas because population controls needed to make estimates of levels comparable to other areas are not available. Combining different datasets for different variables would produce labor force statistics that are not internally consistent as they are based on different survey concepts and samples, and consistency is crucial to derive residual migration from the VAR system.

data for Metropolitan Statistical Areas (MSA), which includes net domestic migration at the MSA level starting in 2005.[20] To construct a measure of predicted relative employment growth (variable $rimix$) at MSA level, we follow equation (3), but using employment shares by industry at MSA level tabulated in the *Geographic Profile of Employment and Unemployment* to weigh the aggregate industry employment growth. We then estimate equation (10) using data for the largest 54 MSA's from 2005-2014 (leaving out lagged dependent variable) and obtain results summarized in column (4) to (6) of Table 3.

As the MSA concept is based on commuting patterns within a metropolitan area, it approximates a local labor market better than do state lines. For example, the tri-state area Maryland-Virginia-District of Columbia partially makes up one local labor market, while the state of Virginia is part of 12 distinct local labor markets as defined by the concept of MSA.[21] Interstate migration thus underestimates the share of population moving across local labor markets (see Molloy et al, 2011). We therefore expect the responsiveness of migration to local shocks to be stronger when estimated across MSA's relative to states, as indeed turns out to be the case in column (4). However, the sample includes wide outliers triggered by migration (and return migration) in the aftermath of Hurricane Katrina.[22] Leaving out observations from 2006, we obtain the result in column (5). In response to a 1 percent relative labor demand shock, as predicted by an MSA's production structure, MSA-level population adjusts through net migration by 0.8 percent 1 year after the shock, double the responsiveness estimated at state-level during roughly the same period. Consistent with previous results, column (7) shows that the differential response of migration during the Great Recession was large and highly statistically significant. Subject to the same relative labor demand shock of, say positive 1 percent, population in an average MSA gained by 2 percent after 1 year

[20] An MSA has at least one urbanized area of 50,000 or more population, plus adjacent territory that has a high degree of social and economic integration with the core as measured by commuting ties.

[21] Migration analysis using MSA data brings its own drawbacks: it does not cover the whole United States, delineations are revised frequently and importantly for our purposes, there are no long time series for MSA-level labor force and employment statistics. The commuting zone concept used for example in Autor et al. (2013) overcomes the first two drawbacks but not the third.

[22] Migration flows in the aftermath of Katrina were among the largest in US history. In our MSA sample, this led to an out-migration rate above 25% in some MSA's in Louisiana in 2005-2006 and in-migration rates above 10% to some MSA's in Texas and California during the same time.

during the Great Recession compared with 0.75 percent during other years. Though higher in absolute magnitudes as expected, this differential effect also represents a 3-fold increase in responsiveness during recessions relative to expansions, similar to the corresponding ratio obtained with state-level estimates.

To sum up, we used different estimation techniques (expanding window regressions, direct business cycle interaction) as well as different data sets (state-level labor force statistics, net migration data) and geographic breakdown (state, MSA) to establish that the response of net migration to local relative shocks increases substantially during recessions.[23] The robustness of this result provides additional support for our identification strategy for the RFIV-VAR model.[24]

5.4 Discussion of mechanisms

We close this section by discussing how the main results relate to the existing literature and point to avenues for future work. The finding that population gains in states experiencing positive relative demand shocks are substantially larger than population losses in states subject to negative relative demand shocks is consistent with Glaeser and Gyourko (2005), who document such asymmetry for decadal population growth in the cross-section of cities experiencing positive versus negative weather "shocks", as well as Notowidigdo (2013), who documents qualitatively similar asymmetry in population response to positive versus negative labor demand shifts that persist over a decade.[25] We contribute to this literature by documenting that such asymmetry is also operative at the business cycle frequency and more importantly, that it has been widening over time.

What could be driving this diverging trend? The increasing lack of ability and/or desire to move

[23]This result is also obtained when we estimate the effect in normal and recessionary periods using a smooth transition approach. See Appendix E for details.

[24]Indeed, it is interesting to note that the BK methodology would not have picked up the increase in migration response during the Great Recession, contradicting the direct evidence using migration and house price data (see Appendix Figure A7).

[25]Glaeser and Gyourko (2005) relate the skewness of city growth rates to the durability of housing, Notowidigdo (2013) relates the asymmetry to a lower incidence of adverse demand shocks as low-skill workers are disproportionately compensated with falling rental prices and public assistance programs. Such asymmetry can also be the result of asymmetric mobility costs across population groups affected by adverse compared to favorable labor demand shocks - groups that may differ by skill but also other characteristics (Bound and Holzer, 2000).

out of adversely affected regions may be because locales experiencing positive demand shocks tend to be those with higher share of skill-intensive industries, which in turn are predominantly concentrated in cities with high and rapidly increasing cost of living (see Moretti, 2013), preventing low-income households from moving in to take advantage of better employment opportunities. Policy changes such as more intensive land use regulation at the municipal level, which disproportionately affects low-income households in high-income cities through rapidly rising rents, could also have been an increasing deterrent to move away from declining regions (Ganong and Shoag, 2015), though evidence in Zabel (2012) does not support the hypothesis that housing supply elasticity matters for net migration response to shocks. Finally, there is consistent micro-data evidence showing that low-income individuals who move to higher-income neighborhoods with more job opportunities do not experience significant improvement in employment rates and earnings, as documented in various studies that follow up on the Moving to Opportunity experiment carried out in the mid-1990's (Katz et al. 2001, Chetty et al. 2015). Relatedly, Yagan (2014) shows that migrants from heavily hit to lightly hit areas during the Great Recession experienced unusually small employment gains. If local disparities in employment access persist at the individual level even after moving to better regions as these studies suggest, then the increasing reluctance of moving away from relatively worse performing states may be at least in part related to these frictions.

However, none of the mechanisms put forward in the literature so far can explain the counter-cyclical migration responsiveness we find here. As the majority of the counter-cyclical responsiveness is accounted for by stronger in-migration to better-performing states, explanations that focus on (possibly time-varying) reluctance to migrate out of worse-performing states mentioned above are unlikely to drive the result. Instead, we think two potential mechanisms could offer promising insights. First, there can be compositional effects. We know that mobility is higher among the unemployed and labor market entrants than the rest of the population.[26] Thus faced with a given differential in employment opportunities across regions, the response of migration into better-performing states would be stronger if the out-of-state population is composed relatively more of

[26]The 1-year interstate migration rate for the unemployed has been 4.74 percent on average in 1976-2013, roughly double the 2.35 percent average interstate migration rate for the overall population.

these mobile groups.[27] As recessions are episodes when the share of unemployed increases across all demographic groups and all regions, this compositional effect could be substantial. To explore possible magnitudes, we use micro data from the March Supplement to the Current Population Survey (CPS) to compute the share of the labor force that moved across state borders to "look for work or lost job", to measure mobility in response to better employment opportunities.[28] Computing the counter-factual job-search migration rate that would result purely due to compositional changes (holding the within group migration rates fixed), we find that the compositional effect only accounts for around 40 percent of the increase in job-search mobility in 2008, while it potentially explains up to 70 percent of the change in 2009 (green line in Figure A10 and last row in Table A3), both relative to the baseline level in 2006. That is, in the initial years of the crisis, the bulk of the increase in interstate migration for job-search is driven by higher migration rates within the groups, particularly those unemployed at least a year and recent labor market entrants.[29] Second, motivated by work such as Crossley and Low (2014), and Hoffmann and Shcherbakova-Stewen (2012), who show that credit constraints are strongest during recessions and most binding for job losers, we believe that counter-cyclical credit constraints may be an important factor pushing people to migrate more during recessions. As acknowledged by BK themselves, liquidity constraints "may force workers who become unemployed to leave the state rather than borrow and wait for the upturn" (BK, p.54). This is also supported by our observation of the data showing that the job-search migration rate increased the most for the long-term unemployed and labor market entrants, who we expect are least able to borrow and hence face the highest consumption risk. We leave a systematic investigation of this mechanism for future research.

[27]This compositional effect may not increase outflows from depressed states as much if the above-mentioned deterrents and frictions persist and work against the increased mobility through composition.

[28]Respondents' stated reason for moving has been included in the March supplement of the CPS starting in 1999. Appendix Figure A10 and Table A3 present the evolution of this job-search-related migration rate, which displays a noticeable increase following the Great Recession consistent with our results using aggregated state-level data.

[29]Changes in the composition of the unemployed and labor market entrants along demographic characteristics (age, skill, occupation) during the Great Recession cannot explain the increase in their respective job-search mobility. In other words, the increase was caused by higher propensity to migrate for job-search among unemployed and new entrants within each demographic group, see Appendix Table A4.

6 Concluding remarks

Over the last 25 years or so, the American population has become less mobile. Mobility has decreased most notably for long-distance movers crossing state borders. Our paper shows that the reduced mobility was also associated with less net migration across states in response to spatial disparities in labor demand conditions, possibly slowing down the diffusion of shocks through the channel of inter-regional population adjustment. Importantly, this has not slowed down the pace at which more productive regions have been attracting in-migration, but has been exclusively driven by a weakening of out-migration from poorer states. Mobility is also less instantaneous than previously thought, placing a much larger burden on the local non-employment pool in the short and medium-run. Notwithstanding the decline from the early 1990s, the response of inter-state migration to regional asymmetries in job opportunities actually increases in recessions, contradicting concerns of increased geographic mismatch often raised in the wake of the Great Recession.

That said, our results leave open the question of whether the trend and cyclical pattern in mobility that we document are efficient and if there is scope for policy to improve welfare by influencing individual migration decisions. To answer this, it is necessary to understand the underlying drivers of these patterns. Our results regarding the asymmetry and counter-cyclicality of population flows between states, combined with the related literature, suggest that migration decisions are determined by various frictions (informational, credit market) and policies (housing, social programs), and that welfare outcomes depend on interactions of these and other interrelated factors. Exploring how frictions and policies can interact in shaping the dynamics of labor mobility remains an important area for future research. More generally, studying spatial patterns of labor market adjustments offers an alternative lens to understanding the workings of the aggregate labor market and the propagation of macroeconomic shocks.

References

Autor, D.H., Dorn, D., Hanson, G.H. (2013), "The China Syndrome: Local Labor Markets Effects of Import Competition in the United States", *American Economic Review* 103(6), pp. 2121-2168

Balakrishnan, R., Dao, M., Sole, J. and Zook, J. (2014), "Recent US labor force participation dynamics: reversible or not?", *IMF Selected Issues Paper* 2014, 5-29, Washington, DC.

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

Beyer, R. C., and Smets, F. (2015), "Has mobility decreased? Reassessing regional labour market adjustments in Europe and the US", *Economic Policy*, October 2015, pp. 643-682.

Blanchard, O., and Katz, L.F. (1992), "Regional evolutions" *Brookings Papers on Economic Activity* 1, 1-77.

Bound, J. and Holzer, H.J. (2000), "Demand shifts, population adjustments, and labor market outcomes during the 1980s", *Journal of Labor Economics* 18(1), 20-54.

Chernozhukov, V., C. Hansen (2008), "The reduced form: A simple approach to inference with weak instruments", *Economics Letters* 100, pp. 68-71.

Chetty, R., Hendren, N., and Katz, L.F. (2015), "The Effects of Exposure to Better neighborhoods in Children: New Evidence from the Moving to Opportunity Experiment", forthcoming *American Economic Review*.

Crossley, T. F., and Low, H. W. (2014), "Job Loss, Credit Constraints, and Consumption Growth", *Review of Economics and Statistics*, 96(5), 876-884.

Decressin, J. and Fatas, A. (1995), "Regional labor market dynamics in Europe", *European Economic Review* 39, 1627-1655.

Gallin, J.H. (2004), "Net migration and state labor market dynamics", *Journal of Labor Economics*, 22(1), 1-21.

Ganong, P. and Shoag, D. (2013),"Why has regional income convergence in the US declined?", Working Paper, December 2013.

Glaeser, E., and J. Gyourko (2005), "Urban Decline and Durable Housing", *Journal of Political Economy* 113(2), pp. 345-374.

Hoffmann, M. and Shcherbakova-Stewen, I. (2011), "Consumption risk sharing over the business cycle: the role of small firms' access to credit markets", *Review of Economics and Statistics*, 93(4), 1403-1416.

Jimeno, J., Bentolila, S. (1998), "Regional unemployment persistence (Spain, 1976-1994)", *Labour Economics* 5, 25-51.

Kaplan, G., Schulhofer-Wohl, S. (2013), “Understanding the long-run decline in interstate migration”, Federal Reserve Bank of Minneapolis Working Paper No. 697, February 2013.

Katz, L.F., Kling, J.R., Lieman, J.B. (2001), “Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment”, *Quarterly Journal of Economics* May 2001, pp. 607-654.

Molloy, R., Smith, C.L. and Wozniak, A. (2011), “Internal migration in the United States”, *Jounal of Economic Perspectives* 25(3), 173-196.

Saks, R. and Wozniak, A. (2011), “Labor Reallocation over the Business Cycle: New Evidence from Internal Migration”, *Journal of Labor Economics* 29(4), 697-739.

Moretti, E. (2013), “Real Wage Inequality”, *American Economic Journal: Applied Economics* 5(1).

Notowidigdo, M. (2013), “The Incidence of Local Labor Demand Shocks”, *NBER Working Paper* 17167.

Partridge, M. D. and Rickman, D. S. (2003), “The waxing and waning of regional economies: The chicken - egg question of jobs versus people”, *Journal of Urban Economics*, 53(1), 76-97.

Partridge, M., Rickman, D., Olfert, M.R., and Ali, K. (2012), “Dwindling U.S. internal migration: Evidence of spatial equilibrium or structural shifts in local labor markets?”, *Regional Science and Urban Economics* 42, pp. 375-388.

Pesaran, M. H., Timmermann, A., (2002), “Market timing and return prediction under model instability”, *Journal of Empirical Finance* 9(5), 495-510.

Roback, J. (1982), “Wages, rents, and the quality of life”, *The Journal of Political Economy*, 1257-1278.

Yagan, D. (2014), “Moving to opportunity? Migratory insurance over the Great Recession”, *Working Paper*, January 2014.

Zabel, J.E. (2012), ”Migration, housing market, and labor market responses to employment shocks”, *Journal of Urban Economics*, 72(2), 267-284.

TABLES & FIGURES

Table 1: Endogeneity of contemporaneous employment growth: Employment rate (le) and Participation rate (lp) equation

	OLS		2SLS	
	Dependent variable:			
	(1)	(2)	(3)	(4)
	le	*lp*	*le*	*lp*
$\Delta e_{s,t}$	0.225***	0.411***	0.792***	0.095
	(0.021)	(0.025)	(0.149)	(0.134)
Hausman p			0.00	0.00
1st stage			*Dep. Var:* $\Delta e_{s,t}$	
$rimix$			0.580***	
			(0.133)	
F-stat.			19.1	
N	1785	1785	1734	

Note: Column (1), (2) show the OLS estimate of le and lp on $\Delta e_{s,t}$ using OLS and and column (3), (4) the 2SLS estimate using the instrument $rimix$ as defined in equations 3 in the text. The first stage panel shows the estimates of the endogenous variable $\Delta e_{s,t}$ on $rimix$. Robust standard errors clustered on states are given in parenthesis. All regressions also include the set of lagged endogenous variables as in each equation of the system in (1) as well as state fixed effects.

Table 2: Test statistics (p-value) for rejecting the null hypothesis of over-identification using Census/IRS migration data.

	OLS-VAR	IV-VAR
Census migration data for m		
t=1	91.14 (0.00)	0.26 (0.61)
t=2	74.26 (0.00)	1.32 (0.25)
t=3	77.88 (0.00)	1.48 (0.22)

Note: The entries present the Chi-squared test statistics (with 1 degree freedom) and, in parenthesis, the p-values for the cross-equation restrictions that correspond to the null hypothesis that the VAR-implied migration response are equal the directly estimated migration response at each of the 3 years after the relative labor demand shock. The first column tests the over-identification in the stacked OLS-VAR model (7) and the second column in the stacked RFIV-VAR model (8) in the text.

Table 3: Direct estimation of counter-cyclical migration response.

	Dependent variable: net migration rate m_t, from:					
	Census PEP			Census PEP (MSA)		
	(1)	(2)	(3)	(4)	(5)	(6)
m_{t-1}	0.417**	0.403***				
	(0.168)	(0.143)				
$rimix_t$	0.041					
	(0.056)					
$rimix_{t-1}$	0.417***			1.349***	0.810**	
	(0.085)			(0.393)	(0.341)	
$D(Exp)_t \times rimix_{t-1}$ (γ_1)		0.272***	0.343***			0.757**
		(0.092)	(0.124)			(0.374)
$D(Rec)_t \times rimix_{t-1}$ (γ_2)		1.013***	1.182***			2.008***
		(0.193)	(0.177)			(0.420)
Time FE	yes	yes	yes	yes	yes	yes
$H_0 : \gamma_1 = \gamma_2$						
p-value		0.01	0.00			0.01
N	1122	1122	1173	447	393	393
Years in sample	1992-2013	1992-2013	1992-2013	2005-2014	dropping 2006 (Katrina)	dropping 2006 (Katrina)

Note: Estimates are based on equation (10) in the text using different sources of net domestic migration data. Census PEP refers to state-level net domestic migration data from the Census Population Estimates Program, ACS migration rates are from the American Community Survey and measured as net inflow of 16-64 year old adults as percent of beginning of year working age population, Census PEP (MSA) refers to MSA-level net migration data from the Census Population Estimates Program. $D(Rec)$ denotes a dummy variable that equals one if the observation year contains at least one NBER recession quarter, and $D(Rec)$ equals one if the observation year contains none. Columns (6) and (7) drop observations in 2006 to exclude outliers from migration waves following Hurricane Katrina. Robust standard errors clustered on states/MSA are provided in parenthesis. All regressions include a set of state/MSA and year fixed effects. Regressions are weighted by state's/MSA's average population over the sample.

Figure 1: Dispersion of Employment Growth Rates across US States, 1977-2015.

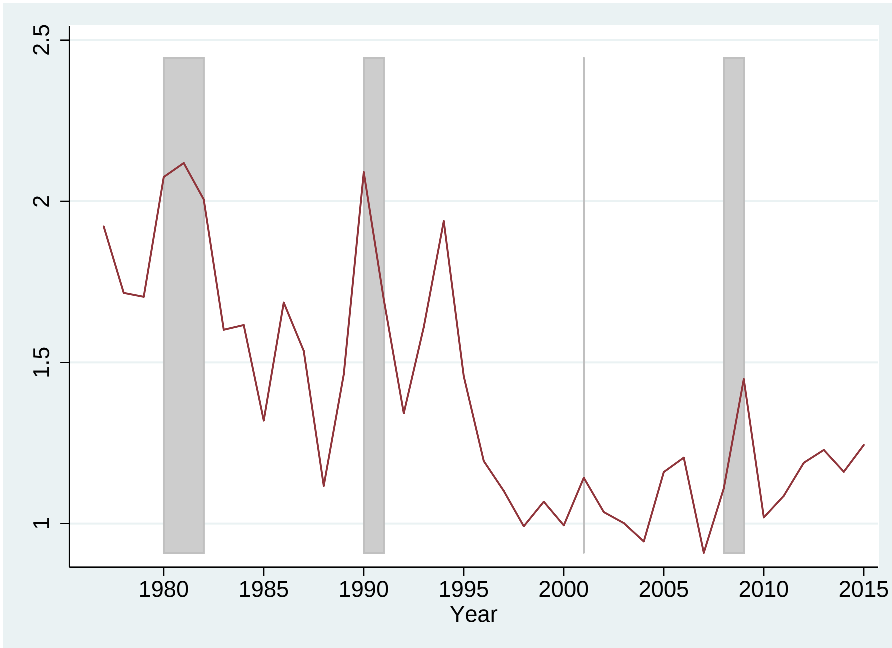

Note: Authors' calculations based on data from the BLS Local Area Unemployment Statistics (LAUS). Each data point corresponds to the standard deviation of employment growth rates across all US States in the given year. Shaded areas represent years with NBER-dated recessions.

Figure 2: Response of state-relative labor market variables: OLS-VAR vs. RFIV-VAR.

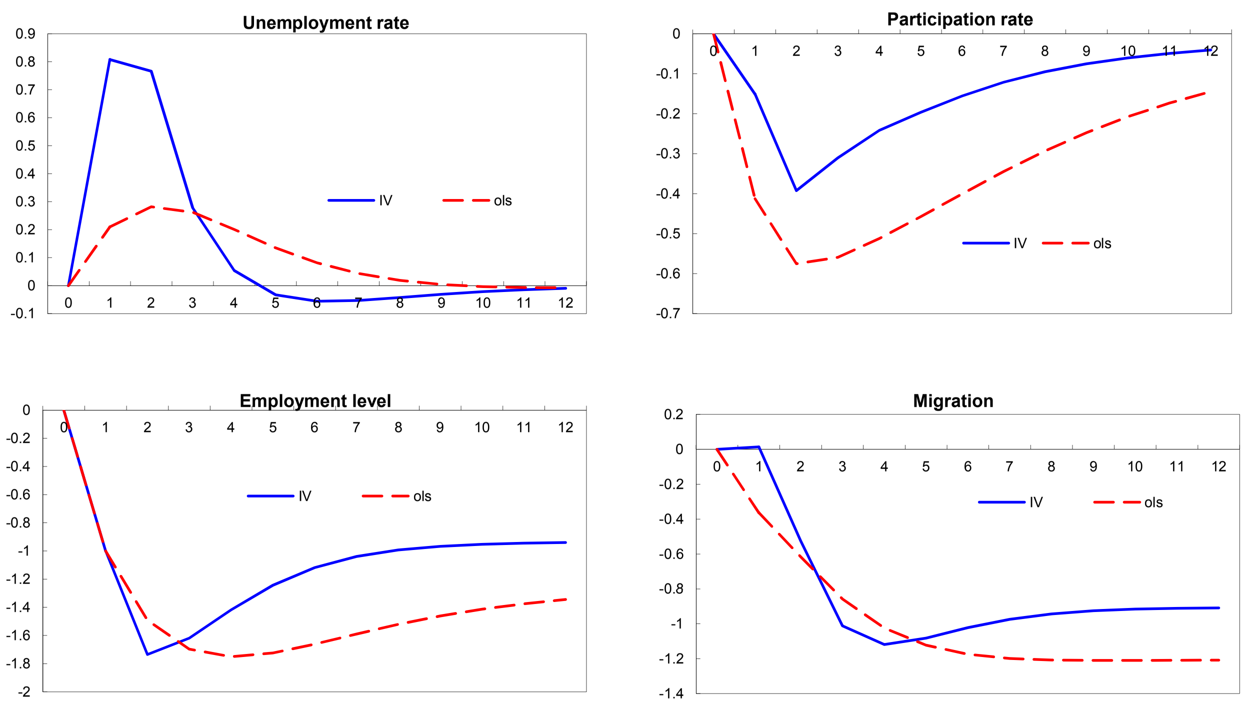

Note: Impulse response to 1 percent relative negative labor demand shock under OLS (derived from equation system 1) and reduced form using $rimix$ as IV (derived from the equation system 4). Units are percent deviation from pre-shock values for employment level, percentage points for unemployment and participation rates, and percent of pre-shock working-age population for net migration.

Figure 3: Response of population, using net migration and population data directly vs. backed-out from VAR.

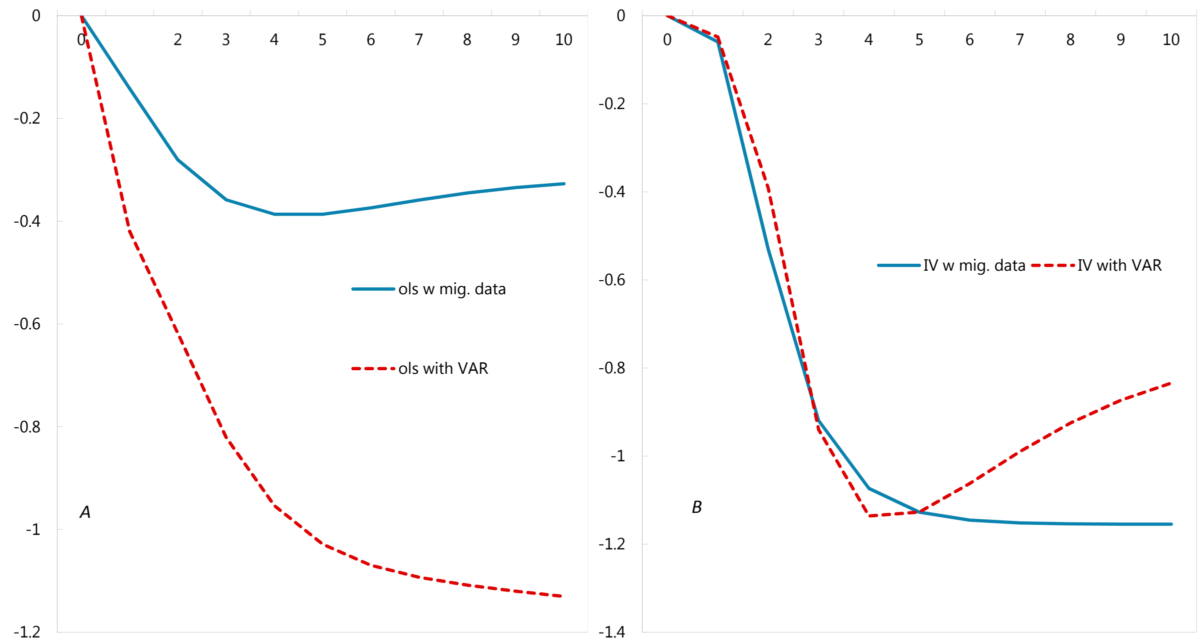

Notes: Sample period is 1991-2013 for net migration data comparison against estimates from OLS-VAR model (1) in the text (A) and RFIV-VAR model (5) in the text (B panel). Horizontal axis denotes years after shock. Unit on vertical axis: percent of working-age population.

Figure 4: Response of population to 1 percent state-relative labor demand shock: sub-samples.

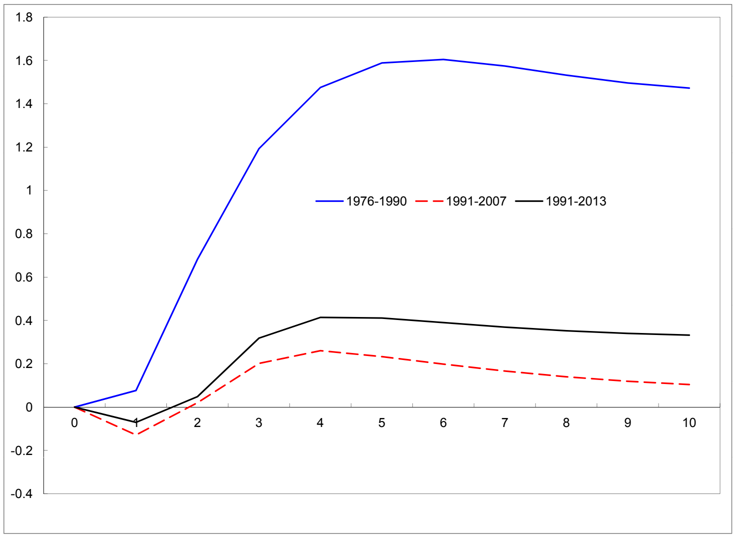

Notes: Each line plots the implied response of state-level population to 1 percent relative labor demand shock ($\Delta rimix = 0.01$) in three sub-samples, estimated with the baseline equation system (4).

Figure 5: Response of population to a 1 percent relative positive vs. negative labor demand shock.

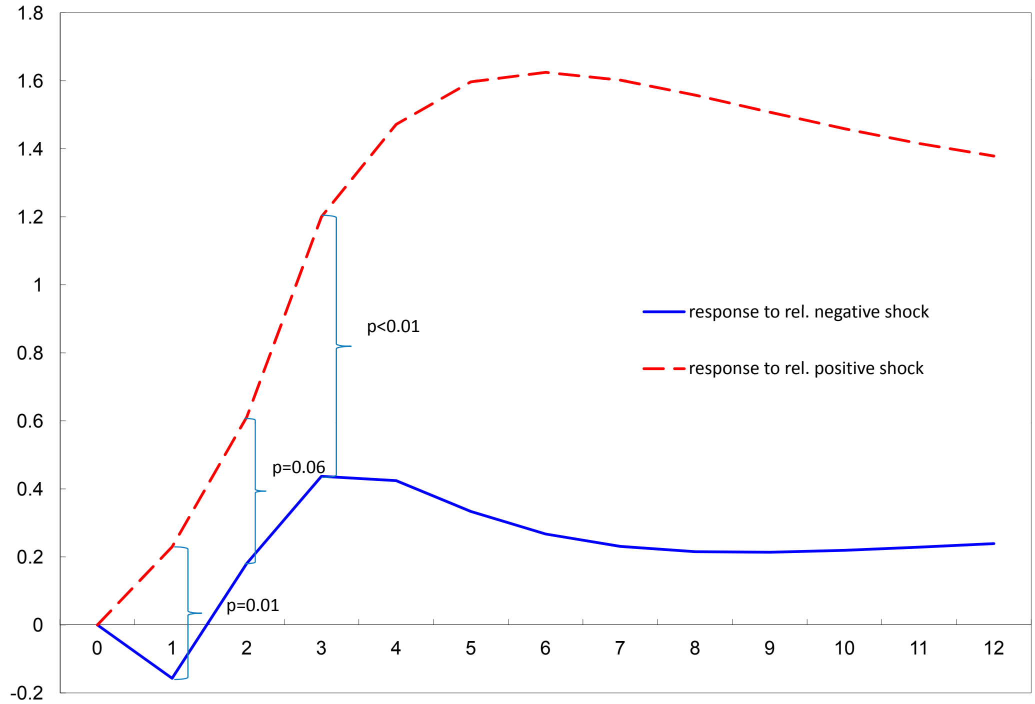

Note: Response of state-level population to a 1 percent relative positive labor demand shock ($\Delta rimix = 0.01$) and relative negative labor demand shock ($\Delta rimix = -0.01$) is derived from the system of equations (9) in the text. P-values for Wald test of equality of responses to positive and negative shocks are derived using the delta method for $t > 0$. Unit on vertical axis is percent of working-age population.

Figure 6: Response of population to a 1 percent relative positive vs. negative labor demand shock: expanding window regression.

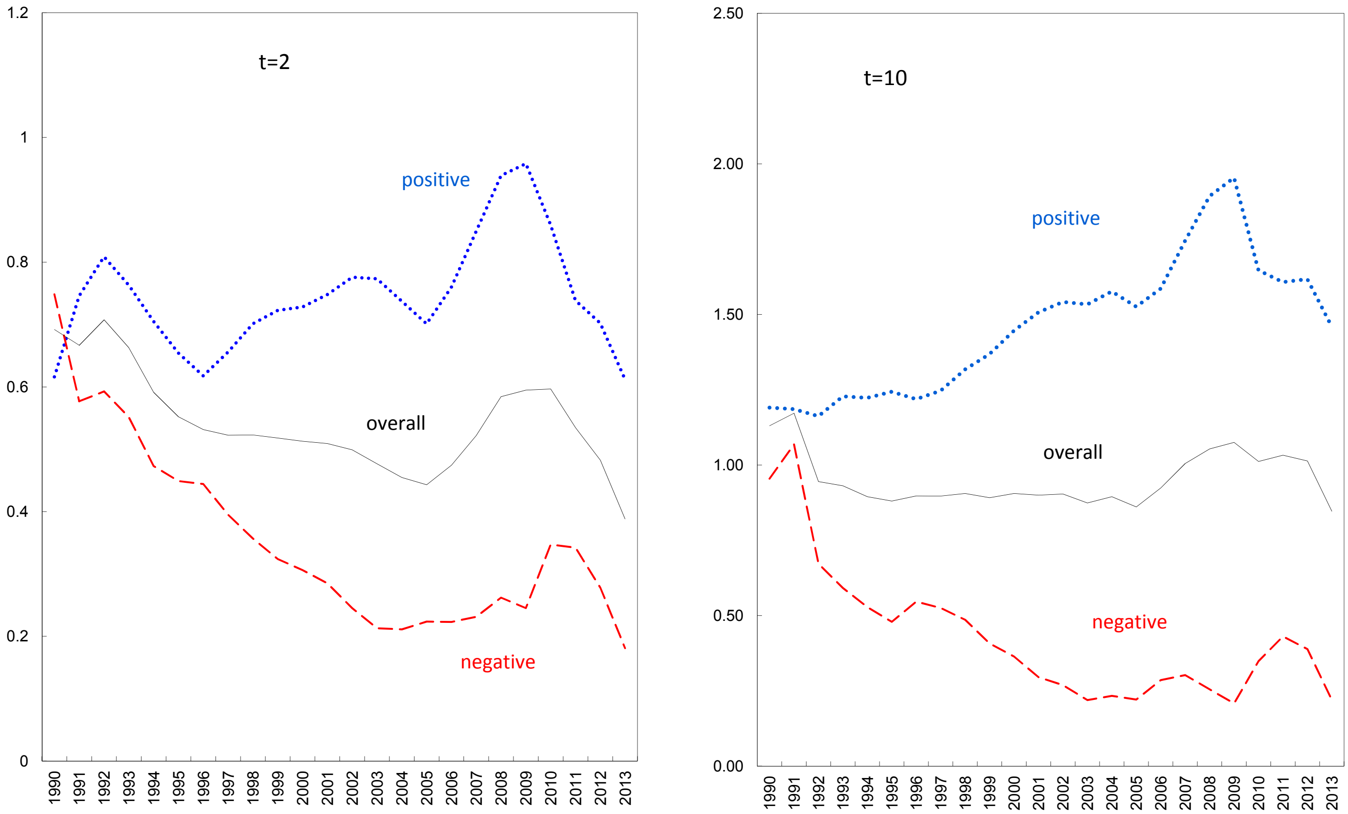

Note: Absolute response of state-level population to a 1 percent relative positive labor demand shock ($\Delta rimix = 0.01$) and relative negative labor demand shock ($\Delta rimix = -0.01$), within 2 and 10 years after the shock, are derived from the system of equations (9) in the text. The overall response is derived from the baseline (symmetric) model (4). Estimates are ordered by last year of observation on the horizontal axis, with the first year being 1977 in all samples. Unit on vertical axis is percent of pre-shock population.