The Persistence of Local Joblessness

Michael Amior and Alan Manning*

7th March 2015

Abstract

Local differences in US employment-population ratios and unemployment rates have persisted over many decades. Using decennial census data from 1950-2010, we investigate the reasons for this. The persistence cannot be explained by permanent differences in amenities, local demographic composition or the propensity of women to work. Population does respond strongly to differences in economic fortunes, although these movements are not large enough to eliminate shocks within a decade. Over the longer run, persistence in local joblessness is largely explained by serial correlation in the demand shocks themselves. The paper also shows that commuting plays an important part in adjustment to local shocks, even when the units of analysis are, as here, commuting zones.

1 Introduction

It is well known that local joblessness persists over many decades. See, for example, Kline and Moretti (2013) and Rappaport (2012) on the US; and Overman and Puga (2002) on Europe. The first panel of Figure 1 compares employment-population ratios (from here on, "employment rates") in 1980 and 2010 among individuals aged 16 to 64, for the 50 largest US commuting zones in 1980. The OLS slope coefficient is 0.5 and highly significant. We argue below that these persistent disparities reflect real gaps in labor market opportunities (for individuals with a fixed set of characteristics): they cannot be explained by permanent differences in local demographic composition. There is also little evidence that they are compensated by local amenities. This is indicative of long-running differentials in welfare.

*Centre for Economic Performance, London School of Economics, Houghton Street, London WC2A 2AE.
We are grateful for feedback from participants at the CEP Annual Conference.

The most natural explanation for the persistence of welfare disparities across areas is constraints on labor mobility, which limit the adjustment to local shocks. But, the second panel of Figure 1 suggests a substantial migratory response. Based on the best-fit line, those commuting zones with the lowest jobless rates in 1980 grew by over 50 percentage points more in the subsequent three decades than those with the highest rates[1], notwithstanding the unusually large growth experienced by many Sunbelt cities (below the 37th parallel). But despite this, the migration responses are insufficient to equalize labor market outcomes over long periods of time.

We argue the missing piece of the puzzle is substantial persistence in local demand shocks. As Figure 2 shows, employment growth between 1950 and 1980 is strongly correlated with growth between 1980 and 2010.[2] Figures 1 and 2 are impressionistic but, in this paper, we derive and estimate a model to understand the causes of persistence in joblessness. The bottom line is that persistence in joblessness is caused by persistence in local demand shocks, together with a migration response that is large but not sufficient to keep up with the demand shocks.

The structure of the paper is as follows. In the next section, we analyze a simple model of the local economy, based on the classic Rosen-Roback framework (Rosen, 1979; Roback, 1982). Unlike much of the literature, we do not assume that local workers supply their labor inelastically. Or alternatively, in the framework of Blanchflower and Oswald (1994), the "wage curve" (the relationship between the local wage and employment rate) is not vertical. This means that local demand shocks affect employment rates and not just wages. Indeed, Blanchard and Katz (1992) and Beaudry et al. (2014b) present evidence that real wages change little along the adjustment path to local demand shocks. If labor supply is not inelastic, we show how one can use the employment rate as a "sufficient statistic" for local economic opportunity - as an alternative to the more common real consumption wage. This approach has precedent in Blanchard and Katz (1992), who implicitly rely on a similar claim in deriving their empirical model. We argue this change in focus from wages to employment rates has practical advantages, as employment rates are easier to measure than real consumption wages for our detailed local geographies. Also, since the employment rate is a stock measure like population, our estimates are directly informative of the speed of population adjustment.

We combine our model of local equilibrium for a fixed population with a simple model of

[1]These population responses were previously documented by Glaeser et al. (1995) and Glaeser and Shapiro (2001).

[2]Similar patterns have previously been documented by Blanchard and Katz (1992) and Dao et al. (2014).

migration: individuals tend to move to areas that offer higher utility, although this process takes time. While our model is set in continuous time, we derive predictions about the change in population between two points in time. This leads to an equation in which the change in log local population is influenced by the change in log employment and the lagged log employment rate: this is essentially a local-level error-correction mechanism (ECM).

To our knowledge, this specification is a new contribution to this topic. Many studies in the recent urban literature have estimated the response of (usually decadal) population growth to *contemporaneous* demand shocks (such as Bound and Holzer, 2000; Glaeser et al., 2006; Notowidigdo, 2011; Autor et al., 2013; Beaudry et al., 2014b; Acemoglu et al., 2014), identified with suitable instruments. But, by excluding the lagged disequilibrium term, there is an implicit assumption that the local economy is in equilibrium at each observation. This jars somewhat with existing evidence of large effects on labor force participation (see e.g. Bound and Holzer, 2000; Autor and Duggan, 2003; Autor et al., 2013), which are suggestive of substantial persistence. Conversely, seminal work by Blanchard and Katz (1992) has focused on the short run dynamics (at annual frequency), estimating a vector-autoregression (VAR) in local employment growth, employment rate and participation rate. But identification in the VAR framework relies on the exclusion of contemporaneous shocks, which is problematic for the longer data frequencies which interest us. The novelty of the ECM specification is to control for *both* contemporaneous employment shocks *and* the lagged disequilibrium term simultaneously.

The third section describes our data, based on decennial census records from 1950-2010 for the 722 Commuting Zones[3] (CZs) of the Continental US. We also report estimates of the persistence in the local employment rate. This is remarkably strong over many decades, for both men and women, and both before and after 1980. Controlling for local demographic composition and time-invariant amenities makes little difference, and the persistence remains strong even after controlling for CZ fixed effects. The persistence is somewhat weaker among higher skilled workers and among labor force participants (as opposed to the general population).

In Section 4, we present estimates of our model for population growth, addressing endogeneity problems using instruments based on industry shift-shares (following Bartik, 1991). The model fits the data well. In our preferred estimates, the elasticity of population to contemporaneous (decadal) employment growth is 0.66, and the elasticity to the initial local

[3]This geographical scheme was originally developed by Tolbert and Sizer (1996) and recently popularized by Autor and Dorn (2013) and Autor et al. (2013).

employment rate is 0.45. This implies a large but incomplete population adjustment over ten years: it corrects for about half the initial deviation in the local employment rate. In specifications controlling for CZ fixed effects or after first differencing (both are demanding specifications in our context), our estimates of population adjustment (and the standard errors) are larger. Evidence of a decline in the size of the migration response over time is mixed, despite a steady fall in regional migration rates since the 1980s (see e.g. Molloy et al., 2011).[4] But, we do find some evidence that high skilled population adjusts more quickly than low skilled.

Although we identify a sizeable population response, it is slower than some previous estimates in the literature suggest. Blanchard and Katz (1992), the seminal study in this field, find the effect of a state-level employment shock on the employment rate disappears after five to seven years.[5] Data frequency may help explain the difference in our results: without access to the data now available, Blanchard and Katz impute long-run responses from an empirical model with two annual lags. This imputation may over-state the long-run pace of adjustment if, for example, it is the most mobile workers who move first. We instead benefit from multiple decades of census data.

The estimates discussed so far condition on the level of employment in an area. But for a complete analysis of the long-run response to a local shock, we also need an equation describing how employment responds to population growth and joblessness. In particular, population growth or slack labor markets may attract new productive investments; and they may also affect the demand for non-traded goods. Following the example of Beaudry et al. (2014b) and Beaudry et al. (2014a), we instrument population growth using climate indicators. Specifically, we use local January temperature. This is motivated by Rappaport's (2007) claim that individuals have increasingly located in cities with mild winters because

[4]Our result may lend support to the view that the declining migration rate is driven by falling *returns* to migration, rather than rising mobility costs (Kaplan and Schulhofer-Wohl, 2012). Having said that, Dao et al. (2014) and Bayer and Juessen (2012) find a clear decline in the population response to local shocks in recent decades.

[5]Dao et al. (2014) and Beyer and Smets (2014) find comparable results after estimating a similar model to Blanchard and Katz (1992) on updated datasets. Using more recent administrative panel data, Yagan (2014) finds that the migration option provided little insurance (in employment terms) to Americans affected by adverse local employment shocks in the Great Recession. But, he argues this is largely due to relatively poor employment outcomes of migrants in their new cities, rather than insufficient migration; so this view does not contradict Blanchard and Katz (1992). On estimating Blanchard-Katz-type models on European regions, Decressin and Fatás (1995), Dao et al. (2014) and Beyer and Smets (2014) find a similar amount of persistence compared to the US. Having said that, other studies find larger persistence in some European countries (see e.g. Hughes and McCormick, 1987; Eichengreen, 1993; Bertola and Ichino, 1995; Obstfeld and Peri, 1998; Beyer and Smets, 2014).

of their amenity value.

We then derive the dynamic response to a local shock, based on our preferred estimates of the population and employment equations. The employment rate is predicted to follow an AR(1) process, with a decadal persistence of 0.55. While this explains most of the decadal persistence in the data, this process clearly cannot match the large persistence after 30 years reported in Figure 1. Instead, we match the data by injecting persistence into the demand shock *itself*: we find that an AR(1) parameter of between 0.6 and 0.8 is needed to achieve this. Of course, we do not directly observe local demand shocks, and these numbers do appear large. But, the persistence in our demand shock instrument (the industry shift-share) falls in the same range; and this reassures us that these numbers are plausible.

Section 6 explores some other outcomes. First, we document the impact of the demand shocks on wages and housing costs, though these are difficult to interpret without a local index of real consumption wages. And second, we consider the role of commuting. We show that changes in commuting patterns *across* CZs play an important role in the adjustment to local shocks. In particular, the share of employed residents who work locally (in their own CZ) is sensitive to how the local area is performing (relative to adjacent CZs). This result is significant, given that labor market outcomes in CZs are usually assumed to be independent in empirical work.

Our main conclusion is that the persistent disparities in local joblessness do indeed reflect persistent differences in labor market opportunities and utility for individuals with fixed characteristics. Migration acts to reduce these disparities, but serial correlation in local demand growth ensures these adjustments are insufficient to equalize utility.

2 A simple model

In this section, we present a slightly modified version of the classic Rosen-Roback model (Rosen, 1979; Roback, 1982) - see also the recent overviews by Glaeser and Gottlieb (2009) and Moretti (2011). There is a single traded good with price P, common across all areas. And there is a single non-traded good, housing, with price P_r^h which does vary across areas, r. Assuming preferences are homothetic (perhaps for convenience more than realism), there is a unique price index in each area given by:

$$E_r = E\left(P, P_r^h\right) \tag{1}$$

On the population side we assume that everyone is homogeneous ex ante, and total population is given by L_r. When in work, we assume individuals earn W_r. In the classic formulation of the Rosen-Roback model, employment and population are synonymous; but it is an important distinction to make in our model, and we denote employment by N_r. We assume (see Albouy, 2012) that labor income is taxed at a rate τ and that those not in work earn a fraction ρ of the local wage (e.g. through unemployment insurance or disability benefits). Total income in an area is then given by:

$$Y_r = (1-\tau) W_r N_r + \rho W_r (L_r - N_r) \tag{2}$$

Using (1) and basic consumer theory, there will be a demand for housing which, in log form, we write as:

$$\begin{aligned} d_r^h &= y_r - p + d^h \left(p_r^h - p\right) \\ &= w_r - p + l_r + \beta (n_r - l_r) + d^h \left(p_r^h - p\right) \end{aligned} \tag{3}$$

where lower case variables denote logs, and the final line substitutes a log-linear approximation of (2).

Now consider the production side of the economy. We assume housing production does not depend on local labor, so we can write the supply curve as[6]:

$$s_r^h = s^h \left(p_r^h - p\right) \tag{4}$$

We also assume production of the traded good does not depend on land. Consequently, the demand for labor in area r can be written as:

$$n_r^d = z_r + n^d (w_r - p) \tag{5}$$

where z_r is an area-specific shifter that will be the source of local shocks. To close the model, we need a labor supply curve or its equivalent in non-competitive models. In the classic formulation of the Rosen-Roback model, labor is assumed to be supplied inelastically so that $n_r^s = l_r$. However, we will assume there is some elasticity in this relationship, so we

[6]We do not consider here the implications of the fact that housing is a durable good and, in part, an irreversible investment. As Glaeser and Gyourko (2005) and Notowidigdo (2011) note, this yields a specification with an asymmetric impact of demand shocks on population, depending on whether the city is growing or shrinking.

work with a formulation:

$$n_r^s = l_r + n^s \left(w_r - e_r\right) + \eta_r \tag{6}$$

where η_r is a region-specific difference in the labor supply curve. This implies there is a positive relationship between the real consumer wage in an area and its employment rate. Although we do distinguish between population and employment, we do not, for ease of exposition, distinguish between labor force and population, so the employment-population ratio and unemployment rate are synonymous (a more complicated model could address this distinction). (6) can be interpreted as an elastic labor supply curve if one thought the labor market was competitive, or as a "wage curve" (Blanchflower and Oswald, 1994) if one thought some other labor market model was more appropriate.

For a given local population, we can solve for the locally-determined variables as a function of the exogenous variables. The workings of this model are well-known but are briefly repeated here. If labor supply is inelastic, an increase in the demand for labor in the area cannot lead to a rise in employment. Consequently, based on (5), the wage in the area must rise. From (3), this raises the demand for housing, so the price of housing in the area must also grow to equilibrate the housing market. The more inelastic the supply of housing, the more the housing price will grow. But, this price rise will not generally offset the growth in wages, so the real consumer wage of workers is higher after the shock.

If we relax the assumption that labor supply is completely inelastic, the effect on wages is more muted as the employment rate also rises. The more elastic is (6) with respect to the wage, the more the variation will show up in employment rates and less in real wages.

We define utility in region r in this model as:

$$u_r = \beta \left(n_r - l_r\right) + \left(w_r - e_r\right) + a_r \tag{7}$$

where a_r is the value of local amenities. Utility is a function of the employment rate and the real consumer wage which, in turn, is a function of local wages and prices. But, it is difficult empirically to disentangle the effect of employment and wages without independent shocks.[7] In particular, a labor demand shock in this model will not yield independent variation in these two variables. In these circumstances, either the employment rate or real consumer

[7]See Beaudry et al. (2012; 2014b; 2014a) for one attempt to do so. They exploit national-level differences in industry-specific employment and wage growth, together with local differences in industrial composition. This gives rise to distinct local industry-weighted wage and employment shift-shares, which have some independent variation.

wage alone can serve as a one-dimensional measure of local labor market conditions: the other variable can be replaced in (7) by using the wage curve (6). In particular, eliminating the real consumer wage gives:

$$u_r = \beta (n_r - l_r) + n^{s-1} (n_r - l_r - \eta_r) + \tilde{a}_r \tag{8}$$

i.e. utility is just a function of the employment rate and the amenity. Blanchard and Katz (1992) implicitly perform a similar transformation in identifying their VAR model in local quantities (employment growth, employment rate and participation rate). One does lose information in doing this; but if there is a single shock, it is not information that could be identified in any case. Note that this transformation does not require us to assume the elasticity of housing supply is the same in all areas, though this will influence the extent to which labor demand shocks affect the employment rate (see Glaeser et al., 2006; Saiz, 2010). Also, note that the area fixed effect $\tilde{a}_r$ in (8) is an amalgam of the amenity effect and any area differences in the "wage curve".

One could perform a similar substitution to eliminate the employment rate from (7) and write everything in terms of the real consumer wage. Much of the literature does this, albeit often implicitly. There is nothing to choose between the approaches theoretically, except in the extreme cases where labor supply is completely inelastic or the real consumer wage completely rigid. Practically though, local wage deflators are available for shorter periods and fewer geographies than employment rates and are notoriously difficult to estimate where they are available at all. The most typical approach is to use only housing costs (though non-traded goods prices are also likely to matter) or a regional price index, the most common of which in the US is the American Chambers of Commerce Researchers Association (ACCRA) index based on the prices of 59 items across 300 cities. But this has been shown to suffer from some serious problems (Koo et al., 2000), the correction of which has large effects on the relative price indices in different cities (Phillips and Daly, 2010). Albouy (2012) describes also how real wage measures are sensitive to the treatment of various taxes and benefits, non-tradable goods expenditure and non-labor income. And aside from measurement concerns, there is one further advantage of the employment rate over the real wage: employment is measured in the same units as population and this allows us to more directly assess the speed of adjustment.

In this section, we have presented a very simple model that forms the basis for the empirical model we estimate below. Appendix A shows how the same type of equation in (8) can be derived from more complicated models with a non-traded goods sector that

employs local labor (as in Moretti, 2010), agglomeration effects (see Glaeser and Gottlieb, 2009; Moretti, 2011 for overviews), endogenous amenities (e.g. Glaeser et al., 2001; Diamond, 2013), frictional labor markets (e.g. Beaudry et al., 2012; 2014b; 2014a), or heterogeneity in skills (e.g. Moretti, 2011; Diamond, 2013).

So far, we have described the short-run equilibrium contingent on a fixed population. We combine this with a simple model for the migration response, in which the population in an area responds to the gap between local and aggregate utility, the latter of which we denote by u. This suggests the following equation:

$$\begin{aligned} \frac{\partial l_r(t)}{\partial t} &= -g\left(u_r(t) - u(t)\right) \\ &= \gamma_0 + \gamma\left(n_r(t) - l_r(t)\right) \end{aligned} \tag{9}$$

where t denotes time and the second line is a linearization of the first. In a steady-state, the "spatial arbitrage" condition of the Rosen-Roback model guarantees that utility is equal in all areas. In this case, a local demand shock will drive up population in an area in response to rising utility; but this will force up house prices, which produces countervailing downward pressure on local real wages and employment rates. If the local demand shock is small enough to have no effect on the aggregate economy, the final equilibrium will have workers no better-off than before, with higher local wages perfectly off-set by higher local house prices. But, we would expect all of this to take time, and (9) embodies a simple adjustment process.

The specification in (9) assumes local population flows depend solely on *current* labor market conditions and amenities, when forward-looking rational agents should also pay attention to expected *future* conditions. If current conditions are a sufficient statistic for future conditions, one can still derive an equation like (9); but the sensitivity to conditions is a mixture of the true sensitivity and the persistence in utility; see Gallin (2004). Gallin also shows how one can derive an Euler equation for the migration decision in the presence of forward-looking agents by including expected future migration as an extra control and instrumenting it. But this approach requires good instruments both for current labor market conditions and future migration, something which is quite demanding. In addition, models with forward-looking agents typically struggle to estimate precisely the discount factor that measures the relative importance of current and future conditions (see, for example, Gallin, 2004; Kennan and Walker, 2011) and often impose a value, thus assuming agents have rational expectations when this is a hypothesis one would like to test. We prefer to stick with

(9), while acknowledging the difficulty of giving a structural interpretation to the estimates.

Given that our data are in discrete time, we next discretize equation (9). Notice it can be written as:

$$\frac{\partial e^{\gamma_n t} l_r(t)}{\partial t} = e^{\gamma t}\gamma_{0r} + \gamma e^{\gamma t} n_r(t) \tag{10}$$

which has as a solution:

$$e^{\gamma_n t} l_r(t) = l_r(0) + \left(e^{\gamma t} - 1\right)\gamma_{0r}/\gamma + \int_0^t \gamma e^{\gamma s} n_r(s)\, ds \tag{11}$$

which can be re-arranged to give:

$$\begin{aligned}[l_r(t) - l_r(0)] &= \int_0^t \gamma e^{\gamma(s-t)} [n_r(s) - n_r(0)]\, ds \\ &\quad + \left(1 - e^{-\gamma t}\right)[n_r(0) - l_r(0)] + \left(1 - e^{-\gamma t}\right)\gamma_{0r}/\gamma\end{aligned} \tag{12}$$

which can be written as:

$$\begin{aligned}[l_r(t) - l_r(0)] &= [n_r(t) - n_r(0)] - \int_0^t e^{\gamma(s-t)}\dot{n}(s)\, ds \\ &\quad + \left(1 - e^{-\gamma t}\right)[n_r(0) - l_r(0)] + \left(1 - e^{-\gamma t}\right)\gamma_{0r}/\gamma\end{aligned} \tag{13}$$

If employment changes at a constant rate over the period, then this can be written as:

$$\begin{aligned}[l_r(t) - l_r(0)] &= \left[1 - \left(\frac{1 - e^{-\gamma t}}{\gamma t}\right)\right][n_r(t) - n_r(0)] \\ &\quad + \left(1 - e^{-\gamma t}\right)[n_r(0) - l_r(0)] + \left(1 - e^{-\gamma t}\right)\gamma_{0r}/\gamma\end{aligned} \tag{14}$$

The above yields the following empirical specification to be estimated on decadal data:

$$\Delta l_{rt} = \beta_0 + \beta_1 \Delta n_{rt} + \beta_2 (n_{rt-1} - l_{rt-1}) + \beta_3' A_{rt} + d_t + \varepsilon_{rt} \tag{15}$$

where t denotes time periods at decadal intervals, Δ is a decadal change, d_t are time effects, A_{rt} is a vector of observable amenity controls, and ε_{rt} represents an unobserved supply-side shock. Notice that equation (15) has the form of an error-correction mechanism (ECM), with local population adjusting to catch up with employment growth Δn_{rt}. The A_{rt} controls

allow the equilibrium employment rate to vary geographically. We also consider specifications where the A_{rt} vector includes time-invariant area r fixed effects. We estimate these specifications in two ways: either by including a full set of area r binary indicators, or by first differencing all variables to eliminate the fixed effects.

This ECM specification is a new model for this literature, which lends itself naturally to an analysis of adjustment speed and persistence. Following the recent urban literature (see e.g. Bound and Holzer, 2000; Glaeser et al., 2006; Notowidigdo, 2011; Autor et al., 2013; Beaudry et al., 2014b; Acemoglu et al., 2014) which studies the population response to long-run (usually decadal) shocks, we control for (suitably instrumented) contemporaneous local employment growth. But, in the tradition of Blanchard and Katz (1992), we also take the dynamics seriously and condition on a lagged disequilibrium term. We next turn to the data we use to estimate (15).

3 Data

3.1 Local population and employment

We identify local labor markets r with the Commuting Zones (CZs) originally developed by Tolbert and Sizer (1996). CZs were recently popularized by Autor and Dorn (2013) and Autor et al. (2013) as an alternative to Metropolitan Statistical Areas (MSAs). MSAs cover only a limited proportion of the US landmass (unlike CZs whose coverage is universal); and there have been changes in MSA definitions over time: this would be particularly problematic for the very long run analysis of this study. Tolbert and Sizer group the full set of US counties into 741 CZs, applying an algorithm to cross-county commuting data from the census of 1990. We restrict our analysis to the 722 CZs on the mainland. Fortunately, county boundaries have been very stable over time, so we can consistently identify these CZs in all census data in our time-frame (since 1950).[8]

Our empirical application is based on the decennial censuses of 1950-2010. The 1940s present distinct empirical challenges[9]; and before 1940, employment status definitions changed

[8]We make just one modification to the Tolbert-Sizer CZ scheme to enable us to construct consistent geographies over time. Specifically, we incorporate La Paz County (AZ) into the same CZ as Yuma County (AZ). Tolbert and Sizler allocated La Paz and Yuma to different CZs, but the two counties only separated in 1983.

[9]Clearly, the outbreak of war creates an important structural break. And also, the process governing local employment changes was markedly different in the 1940s as a consequence of New Deal public works programs at the beginning of the decade.

on a regular basis.[10] Our main estimates are based on local population and employment counts for all individuals aged 16-64. But, we also estimate the model separately for different education, gender and age groups. We describe how we construct the data in Appendix B below.

As we explain in the following section, OLS estimates of (15) are likely to be biased; so credible identification requires an instrument. In keeping with much of the literature [11], we rely on the industry shift-share variables b_{rt} originally proposed by Bartik (1991). The idea is to assume that, over a decade, the stock of employment in each industry i grows at the same rate in every area r, where this growth rate is estimated using national-level data. Specifically:

$$b_{rt} = \sum_i \phi^i_{rt-1} \left(n_{i(-r)t} - n_{i(-r)t-1} \right) \tag{16}$$

where ϕ^i_{rt-1} is the share of workers in area r at time $t-1$ employed in industry i. The term $n_{i(-r)t} - n_{i(-r)t-1}$, expressed in logs, is the growth of employment nationally in industry i, *excluding* area r. This modification to standard practice was proposed by Autor and Duggan (2003) to address concerns about endogeneity to local employment counts. We use the contemporaneous Bartik shift-share b^N_{rt} as an instrument for current employment growth Δn_{rt}, and we use the lagged shift-share b_{rt-1} to instrument for the lagged employment rate $(n_{rt-1} - l_{rt-1})$. The intuition for the latter instrument is that the employment rate at any point in time could be written as a distributed lag of past labor demand shocks. In practice, it is sufficient to instrument using the first lag alone. We construct the instrument using 2-digit industry data from the IPUMS micro-data: see Appendix B for further details.

3.2 Amenity controls

We take a number of approaches to controlling for the amenity in (15). One approach is to model the amenity as an area fixed effect. Another is to include various controls to measure the amenity directly. Because some of these controls might be expected to have varying

[10] In particular, individuals were defined as employed before 1940 if they were at work on the reference *day*; but since 1940, they were only required to have worked one hour in the reference *week*. This change in definition is likely to have had divergent effects on employment counts in different industries. See the IPUMS webpage for further details: *http://usa.ipums.org/usa-action/variables/EMPSTAT#comparability_section*. Also, the 1920 census was conducted in January (as opposed to April since 1930), and this has severe implications for comparability in agricultural areas (given the seasonality of employment).

[11] See, for example, Blanchard and Katz (1992); Bound and Holzer (2000); Notowidigdo (2011); **?**; Beaudry et al. (2012; 2014b; 2014a).

effects over time, we interact each of them with a full set of year effects in the regressions below.

First, we use county-level data kindly shared by Jordan Rappaport on the presence of coastline (ocean or Great Lakes). Coastline may provide consumption or productive amenities (Rappaport and Sachs, 2003) or physical constraints on population expansion (Saiz, 2010).

Second, we control for some climate indicators. In the sample period, population has grown disproportionately in the Sun Belt region. A popular explanation is the advent of air conditioning, which facilitated a more comfortable and productive life in hotter regions (Oi, 1996). But, Rappaport (2007) shows that Americans have increasingly located in cities with *pleasant* weather, specifically cool summers with low humidity and warm winters. And he argues that a growing valuation of climate amenities can help explain observed trends in Southern population, driven perhaps by rising incomes.[12] However, based on trends in local prices and wages, Glaeser and Tobio (2007) argue that climate is not the primary cause. They propose that, before 1980, population was drawn to the South largely because of improving productivity; and after 1980, the key factor was an elastic supply of housing. Following Rappaport (2007), we control for the maximum January temperature, maximum July temperature and mean July relative humidity. We take mean county-level data on temperature from the Center for Disease Control and Prevention[13] and on relative humidity from the United States Department of Agriculture[14]. We aggregate county-level climate data to CZ-level using land area weights.

Third, we control for log population density in 1900. This measure can proxy for the pull of under-developed land or "frontier" regions. Alternatively, there may be consumption or productive amenities (or disamenities) associated with population density. We use a historical measure of density (measured in 1900) to ease concerns over endogeneity.[15]

We also control for an index of CZ isolation. Specifically, this is the log distance to the closest CZ, where distance is measured between population-weighted centroids in 1990.[16]

[12] In particular, Rappaport finds that hot humid summers have deterred population growth, *controlling for winter temperature*. This is inconsistent with an important role for air conditioning.

[13] Based on the period 1979-2011. See *http://wonder.cdc.gov/*

[14] Based on the period 1941-70, from the Natural Amenities Scale study by McGranahan (1999).

[15] These densities are estimated using county-level population and area data from NHGIS. There have been some changes in county boundaries in the intervening period, and we impute CZ-level data using land area allocations based on shapefiles made available by NHGIS.

[16] Population-weighted centroids for counties in 1990 are estimated by the Missouri Census Data Center: *http://mcdc.missouri.edu/websas/geocorr90.shtml*. We estimate CZ centroids by computing the population-weighted averages across the latitudes and longitudes of county centroids.

Isolation may matter for two reasons. First, it might be considered an amenity or disamenity. And second, it limits opportunities for cross-CZ commuting. Indeed, we emphasize in this study that CZs are not distinct local labor markets: many workers commute across CZ boundaries. As noted above, this isolation measure (like other controls) is interacted with a full set of year effects in the regressions below.

Finally, an important contribution to local population growth arises from foreign migration. Of course, local inflows of foreign migrants are partly a response to local employment growth. But, as is well known, migrants are often guided in their location choice by the "amenity" of established co-patriot communities.[17] In the empirical migration literature, there has been a long tradition (popularized by Card, 2001) of proxying these preferences with historical local settlement patterns. Following Card, we construct a "shift-share" predictor m_{ct} for the contribution of foreign migration to local population growth:

$$m_{ct} = \frac{\sum_o \phi^o_{rt-1} M^{new}_{o(-r)t}}{L_{rt-1}} \tag{17}$$

where ϕ^o_{rt-1} is the share of population in area r at time $t-1$ which is native to origin o. $M^{new}_{o(-r)t}$ is the number of new migrants arriving in the US (excluding area r) between $t-1$ and t. The numerator of equation (17) then gives the predicted inflow of all migrants over those ten years to area r. This is scaled by L_{rt-1}, the initial population of area r. Similarly to the Bartik industry shift-shares above, the exclusion of area r helps allay concerns over endogeneity of the shift-share measure to the dependent variable, local population growth Δl_{rt}. We construct this migrant shift-share variables using census and ACS micro-data from IPUMS, based on 79 origin countries: see Appendix B for further details.

3.3 Estimates of the persistence of the local employment rate

Before moving to the estimates of the population response, we present some estimates of our main object of interest: the persistence of local joblessness. In the first row of Table 1, we set out the autocorrelation function (ACF) of the time-demeaned log local employment rate in the data.[18] The persistence is very strong: the ACF remains above one half even at the sixth (decadal) lag. This reflects the patterns of Figure 1 above. The persistence is similar before and after 1980 (see rows 2 and 3), but weaker among the labor force (row 4) compared to the broader population. And it is largely driven by low skilled workers (see rows 5 and

[17] For example, because of job networks (Munshi, 2003) or cultural amenities (Gonzalez, 1998).

[18] For each decadal lag, we estimate the correlation coefficient using a pooled sample of all CZs and all census years with available data.

6), especially at the larger lags.

There are several supply-side explanations for this persistence, such as local variation in preferences for work, demographic composition, or compensating amenities. These can theoretically give rise to long-run differences in steady-state local employment rates. But, the evidence suggests these theories cannot explain much of the persistence we observe.

For example, the persistence might be driven by geographical variation in women's preferences for labor market participation. Indeed, the ACF at smaller lags is larger for women then men (see rows 7 and 8), but the difference is not large: the correlation coefficients are 0.90 and 0.74 respectively at the first lag. Alternatively, local variation in demographic composition may be responsible. For example, skill shares vary across cities; and low skilled workers face lower employment rates. But, adjusting local employment rates for composition makes little differences to the result (see row 9).[19] A third hypothesis is permanent differences in amenities which compensate individuals for labor market conditions. We test this by purging employment rates of the climate, coastline and population density controls, as well as the migrant shift-share variables (which proxy the attractiveness of CZs for new migrants).[20] But again, this makes little difference to the observed persistence (row 10).

Certainly, unobserved local variation in amenities or demographic variation may play an important role - though we believe we have controlled for the most important effects. As it turns out, even controlling for a full set of state fixed effects[21] has little effect on the ACF (row 11), at least at the first few lags. It is also noteworthy that, at these smaller lags, the level of persistence is similar both within and between states. The between-state results are reported in row 12, where we trace out the ACF for state-level employment rates.

[19]We purge local employment rates of observable demographic characteristics in the following way. For each cross-section of the census/ACS micro-data, we run a logit regression of employment on a range of characteristics (age and age squared; four education indicators, each interacted with age and age squared; a gender dummy, interacted with all the earlier-mentioned variables; and black, Hispanic and foreign-born indicators) and a set of location fixed effects (where "locations" are the finest geographical indicator available in the census micro-data. Based on these estimates, we then predict the average employment rate in each location - assuming the local demographic composition in each location is identical to the national composition. We then estimate CZ-level data by weighting the location data by appropriate population allocations (see Appendix B). Unfortunately, we do not have sub-state location indicators in 1960, so we exclude that year from our data series.

[20]Specifically, we regress the log employment rate on three climate variables (the maximum January and July temperatures and mean July relative humidity), a dummy for the presence of coastline, the log population density in 1900, the log distance to the closest CZ centroid, and the immigrant shift-share, with observations weighted by the local population share. All the above controls, excluding the immigrant shift-share, are also interacted with a full set of year effects. The purged employment rate observations are the residuals from this regression.

[21]Some CZs do straddle state boundaries, so we allocate these to the state accounting for the largest population share of the CZ.

In the final rows, we study the effect of including CZ fixed effects. Given that the panel is short (only seven observations per CZ), the local means are estimated with substantial error. This causes a downward bias in the ACF, as removing the fixed effects introduces an artificial negative correlation between the employment rate and its lag (as the panel is short). This is clear from the implausibly negative correlation coefficients at the larger lags in row 13. Fortunately, this bias is quantifiable. Appendix C shows how one can derive bias-corrected estimates of the true ACF from the sample ACF, although an identifying assumption is required. Our approach is to fix the ratio π of the sixth to fifth autocorrelation: we report results for $\pi = 0.9$, $\pi = 0.5$, and $\pi = 0$. The ACF for $\pi = 0.9$ looks similar to the the basic ACF in row 1. But, the decay is larger for smaller π: the correlation for $\pi = 0$ is strong at the first few lags, though the fixed effects explain much of the correlation thereafter. We argue below that these fixed effects are likely to be picking up persistence in local demand growth. Broadly speaking, the message of Table 1 is one of strong persistence in local joblessness over long periods of time.

4 Estimates of population response

4.1 Basic estimates

In this section, we report estimates of (15). Panel A of Table 2 reports OLS and IV estimates for three specifications: (1) a "basic" specification, which includes the amenity controls, interacted with year effects; (2) a fixed effect specification, in which the time-invariant part of the amenity is modeled as an area fixed effect; and (3) a first-differenced specification (another way to eliminate a time-invariant fixed effect), where the dependent variable is the double differenced log population, and the endogenous right-hand side variables are the double differenced log employment and the lagged change in the log employment rate. We weight all observations by local population share, and standard errors are clustered by CZ.

It is worth emphasizing that the fixed effect and first differenced specifications are empirically very demanding, and other studies in this literature (such as Blanchard and Katz, 1992; Autor et al., 2013; Beaudry et al., 2014b) have not estimated them. In particular, given the short time dimension (just six observations), it is difficult to empirically disentangle a "genuine" amenity fixed effect from long run persistence in the employment rate driven by sluggish migratory adjustment. We might then expect the fixed effect and first differenced specifications to overstate the migratory response to initial local deviations in the employment rate.

In all three OLS specifications, the signs on the variables are what we would expect. Higher contemporaneous employment growth is associated with higher population growth; and the estimated coefficient does not vary greatly across specifications, falling at around 0.8. The coefficient on the lagged employment rate is positive, implying that areas that are doing better initially tend to gain population in the subsequent decade. This effect is much more sensitive to specification, ranging from 0.2 to 1.

But, there are good reasons to think that the OLS estimates suffer from a range of biases. First, population and employment growth are clearly jointly determined[22], so the coefficient on employment growth cannot be interpreted as causal; this should bias the OLS estimate of β_1 upwards.

Second, if our controls do not absorb all the local variation in amenity value, the omitted part of the amenity is likely to be negatively correlated with the employment rate. This is because, in spatial equilibrium, individuals trade off local labor market conditions with compensating amenities. If more valuable amenities are associated with faster population growth in our sample, this will then bias the OLS estimate of β_2 downwards. To the extent that the amenity component is time-invariant, this bias is addressed in the fixed effect specification. But the fixed effect specification is likely to suffer from a "Nickell bias" (Nickell, 1981): de-meaning introduces an artificial correlation between the lagged employment rate (which contains a lagged dependent variable: population) and the regression error.

To deal with these problems, we instrument both the variables of interest in the estimation of (15). Specifically, we use current and lagged Bartik shocks as described in the data section above. As we note in Section 2 above, the transmission from these shocks to employment (in the first stage) will be influenced by local differences in the elasticity of housing supply (Glaeser and Gyourko, 2005; Glaeser et al., 2006; Saiz, 2010; Notowidigdo, 2011). But, we find our first stage regressions have sufficient power without accounting for this heterogeneity.

We would expect the current Bartik shock to be most strongly correlated with current employment growth and the lagged Bartik shock to be most strongly correlated with the lagged employment rate. The first stages are reported in Panel B of Table 2. In the basic specification, this pattern materializes very strongly. The instruments have considerable power and, remarkably, sufficient independent variation despite substantial serial correlation in the Bartik shock (see Figure 5 below and the discussion that follows). The current Bartik shock has a strong effect on contemporaneous employment growth and no effect on the lagged employment rate; and vice versa for the lagged Bartik shock.

[22]Indeed, we estimate the employment response to population in Section 5 below.

But, when we introduce fixed effects or estimate in first-differences, the patterns of correlations between the endogenous variables and the instruments become more complicated (though the instruments remain powerful). We prefer the basic IV specification, despite the fact that the p-value (not reported here) for joint significance of the fixed effects is negligible. This is for two reasons. First, the first stage results in the basic specification are more compelling; and second, as noted above, it is difficult empirically to disentangle a fixed effect from long run persistence in a short panel.

The IV estimates are reported in the second part of Panel A of Table 2. The IV estimate of β_1, the elasticity to contemporaneous employment growth, is 0.66 in the basic specification. As expected, this falls below the OLS estimate, though the reverse is true under fixed effects. The IV coefficient on the lagged employment rate, β_2, is 0.45 in the basic specification, larger than the equivalent OLS estimate (this is also true in the fixed effect specification, but not in first differences).

One might be concerned that the IV results are driven by a few outliers, but the graphical analysis in Appendix D shows that this is not the case. We also show in the appendix that the results are robust to the exclusion of population weights: weighting makes little difference to the basic specification, and the unweighted fixed effects and first differenced coefficients are smaller and closer to the basic specification. Also, the amenity controls make little difference to our β_1 estimate, though omitting the climate controls does yield smaller population responses to the lagged employment rate.

As we emphasize in the introduction, a novel feature of our specification is the inclusion of the lagged error-correction term. It is strongly significant in all specifications in Table 2. We show in Appendix D that omitting the error-correction term (following the example of other studies in the urban literature, as outlined in the introduction) yields a larger coefficient on contemporaneous local employment growth: 0.82 compared to the 0.66 in the basic specification of column 4. This is to be expected, given the serial correlation in the Bartik instrument (see Figure 5 below).

Broadly speaking, the results in Table 2 show there is a robust relationship between population growth on the one hand and employment growth and the lagged employment rate on the other. One way to summarize the results is to consider what the estimates imply about the impact of a change in employment, starting from a position of equilibrium (for reasons we explain later, this is not the same as the impact of a shock to labor demand). Based on the basic IV estimates (in column 4), a 10% fall in the level of employment causes population in the first decade to fall by 6.6%, a large change but not enough to remove the

impact of the shock within a decade. As the employment rate is then 3.4% lower after 10 years, population continues to decline by 1.5% in the second decade (according to our β_2 estimate). So after 20 years, the employment rate would be 1.9% lower than the pre-shock level.

Our model in Section 2 does not distinguish between the working-age population and the labor force: that is, it does not consider the participation decision. Table 3 reports estimates of our basic equation when we substitute labor force for population. Specifically, the dependent variable becomes the change in log labor force; and the lagged employment rate is now measured relative to the labor force rather than population. Similar patterns emerge, but the labor force responds more strongly to both regressors than population. In particular, the β_2 estimates are insignificantly different to 1. This is in line with the conclusions of Beaudry et al. (2014b), though they do not study the inactive population. Taken together, Tables 2 and 3 suggest that any sluggishness in the population response to initial local employment rate differentials is associated with changes on the participation margin. The importance of the participation margin tallies with findings from Autor and Duggan (2003),Rappaport (2012) and Autor et al. (2013), who identify large declines in economic activity (and rising take-up of disability benefits) in cities subject to adverse shocks. It should not come as a surprise then that much of the persistence in local joblessness is driven by inactivity, rather than unemployment (as demonstrated in Table 1).

Our results do imply a sizeable migration response, but smaller than that suggested by Blanchard and Katz (1992). Blanchard and Katz estimate a state-level VAR (in local employment growth, employment rate and participation rate) with two lags, using annual data from the Current Population Survey (CPS) over the period 1978-1990. At the time, this was probably the best available data; but this is a very different frequency from ours and may explain the difference in results. For example, in a comment attached to Obstfeld and Peri (1998), Antonio Fatas raises concerns about long run predictions from a model with just two lags at an annual frequency. One particular issue is that the most mobile workers are the quickest to respond to an employment shock, so projecting the initial responses forward may bias upwards the long run population response.

Recently, there have been concerns about a decline in gross migration rates since 1980, and whether this means migratory responses to local shocks have become more sluggish (see e.g. Molloy et al., 2011). Table 4 reports estimate of our population growth equation for two sub-samples: before and after 1980. The fixed effects and first difference specifications do suggest larger population responses (to both variables of interest) in the early period,

though these specifications are especially demanding given a time dimension of just four observations. In contrast, the basic specification gives little support for this hypothesis: while the coefficient on employment growth is larger in the early period, the coefficient on the lagged employment rate is somewhat smaller. This is consistent with the fact that the persistence in the employment rate is similar both before and after 1980 (see Table 1). Having said that, using a methodology akin to Blanchard and Katz (1992), Dao et al. (2014) and Beyer and Smets (2014) do find evidence of smaller population responses in recent decades.

4.2 Demographic decomposition

The estimates above are based on local aggregates of employment and population stocks, across all working-age individuals. But, these estimates may mask divergent patterns within more detailed demographic groups. In particular, it has often been argued that migratory flows of the lower skilled workers are less sensitive to economic condition. We address these concerns by estimating models separately for different demographic groups. To this end, we use group-specific stocks for all relevant variables: the population growth outcome, and the employment growth and lagged employment rate regressors. We also construct our migrant shift-share controls for specific demographic groups (see Appendix B). But we use the same two instruments as before: the current and lagged Bartik shift-shares.

Table 5 reports estimates for college graduates and non-graduates. In line with other work (e.g. Bound and Holzer, 2000; Wozniak, 2010; Notowidigdo, 2011), we find some evidence of larger population responses among higher skilled workers, at least in our basic specification[23]. This is also consistent with our finding in Table 1 that local employment rates are more persistent among lower skilled labor. In all specifications, the response to contemporaneous employment growth is larger for college graduates, though the gap varies. In our (preferred) basic IV specification, β_1 takes 0.87 for graduates and 0.64 for non-graduates. In the same specification, the elasticity β_2 to the lagged employment rate takes 0.73 for graduates and 0.45 for non-graduates; though it is larger among non-graduate workers in the fixed effect and first differenced specifications.

[23]There has been some debate on why lower skilled workers are less likely to leave cities suffering slumps in demand; see also Topel (1986) and Wozniak (2010). The low skilled are commonly thought to face prohibitive migration costs; and Kennan (2013) provides structural evidence for this hypothesis. Notowidigdo (2011) argues they are relatively sheltered from local demand shocks because of declining housing costs (low income families spend a larger share of their income on housing) and transfer payments. Amior (2014) claims the obstacle to low skilled mobility is meager returns to employment, amplified by limited search effort (by both firms and workers) and job creation.

Table 6 presents IV estimates by gender and for three broad age groups. Again, each regression uses the same two Bartik instruments. The response to current employment growth is somewhat larger for men and 25-44s in all specifications. But, in the basic specification, our β_2 estimate is largest among 45-64s. We report the first stages of all these estimates in Appendix E.

Overall, the evidence suggests our simple model for population growth performs well and is a useful framework for thinking about the impact of demand shocks on employment rates.

5 Dynamic response to demand shocks

5.1 Estimates of employment response

The estimates we have presented allow us to make predictions about how population responds, conditional on the level of employment in an area. But this is is not the same as the response to a shock to labor demand, because employment itself is endogenous. In the simple model described earlier, employment is endogenous purely through an effect on the real producer wage. And if the adjustment of labour demand is sluggish (like population), we can derive a similar ECM equation for employment growth:

$$\Delta n_{rt} = \gamma_0 + \gamma_1 \Delta l_{rt} + \gamma_2 \left(n_{rt-1} - l_{rt-1}\right) + \gamma_3 b_r + \varepsilon_{rt} \tag{18}$$

This is identical to the population response equation, but with population and employment reversed. In reality, labour demand is likely to respond to population growth directly - and not just through wage changes. This is because an expanding population should trigger rising demand for non-traded goods (see, for example, Moretti, 2010). Also, slack labour markets may encourage firms to locate to area r, even if low wages do not. These intuitions are consistent with the specification of (18), but will affect its interpretation.

To estimate this equation, we need instruments for both Δl_{rt} and $(n_{rt-1} - l_{rt-1})$. Naturally, for the latter, we use the lagged Bartik shift-share, while controlling for the current shift-share in the regression. But, it is harder to identify a convincing instrument for population growth, which is exogenous of changes in labor demand. We do have a range of amenity variables, but it is not obvious whether these satisfy the exclusion restriction.

There is precedent in the literature for estimating this kind of equation. Beaudry et al. (2014b; 2014a) study the impact of changes in the log labor force (rather than population) on changes in log employment and the log employment/labor force ratio. They identify this

elasticity using a range of climate indicators: January and July temperatures and rainfall. In the first cited paper, they also use a migrant shift-share instrument, though we find this variable has little explanatory power for local population[24].

Instead, we turn to the climate controls. Hot and humid summers may conceivably be linked to demand growth, if the expansion of air conditioning during the sample period increased labor productivity. Our strategy is therefore to control for July heat and humidity in the regression, and use winter January temperature as an instrument. This follows Rappaport's (2007) reasoning that, as incomes have grown, Americans have been drawn to cities with mild winters (which have value as a consumption amenity). Since the effect of January temperature may change over time, we also include the interaction of January temperature and a time trend as a further instrument.

The results are reported in Table 7. Across all specifications, the OLS elasticity of employment with respect to population is close to 1. This is larger than the IV elasticity from the basic specification (in column 4) of 0.69. This is to be expected: we know from Table 2 that population responds positively to employment. In comparison, Beaudry et al. (2014b; 2014a) estimate an IV elasticity of employment with respect to the labor force of close to 1. The difference in our results is largely explained by the different variables we use. In results not reported here, we estimate a larger coefficient of 0.9 when we substitute labor force for population in equation (18).

Of course, January temperatures may be statistically conflated with other relevant factors. In particular, Glaeser and Tobio (2007) argue that population growth in the Sun Belt was largely driven by growing productivity before 1980 and lax housing regulation thereafter. As a further test (in results not reported here), we attempt to pick up some of these effects by controlling for a full set of interactions between the nine census divisions and years. As it happens, our estimates are little affected: the impact of population growth in the basic IV specification changes from 0.69 to 0.68. Having said that, we remain less confident about these instruments than the Bartik shift-shares we use in the population equation.

Next consider γ_2, the coefficient on $(n_{rt-1} - l_{rt-1})$. In our basic IV specification, a 1% deviation in the initial employment rate leads to a -0.32% decrease in subsequent employment growth. Again, this seems sensible: tighter labor markets (with higher wages) should presumably discourage job creation.

The IV estimator with fixed effects has little power, and the first differenced results

[24]See column 7 (our preferred specification) of Table 12 below. This suggests there may be geographical displacement of native-born individuals.

appear unrealistic. But, the identification (purely through the interaction between January temperature and the time trend) is clearly demanding. Given these concerns, we rely on the estimates from the basic IV specification (in column 4) in the analysis which follows.

5.2 Impulse response functions

Our estimates from the population and employment equations (15 and 18) allow us to simulate the dynamic response to a local labor demand shock. Specifically, we rely on the following simultaneous equation model:

$$\Delta n_{rt} = \gamma_0 + \gamma_1 \Delta l_{rt} + \gamma_2 \left(n_{rt-1} - l_{rt-1} \right) + z_{rt} \tag{19}$$

and

$$\Delta l_{rt} = \beta_0 + \beta_1 \Delta n_{rt} + \beta_2 \left(n_{rt-1} - l_{rt-1} \right) \tag{20}$$

where z_{rt} represents the demand shock. From these two equations, one can derive a model for the evolution of the employment rate:

$$x_{rt} = \left[1 - \frac{\left(1-\gamma_1\right)\beta_2 - \left(1-\beta_1\right)\gamma_2}{1-\beta_1\gamma_1} \right] x_{rt-1} + \left[1 - \frac{\beta_1\left(1-\gamma_1\right)}{1-\beta_1\gamma_1} \right] z_{rt} \tag{21}$$

where

$$x_{rt} = n_{rt} - l_{rt} - \frac{\left(1-\beta_1\right)\gamma_0 - \left(1-\gamma_1\right)\beta_0}{\left(1-\gamma_1\right)\beta_2 - \left(1-\beta_1\right)\gamma_2} \tag{22}$$

is the deviation of the log employment rate from its steady-state. Notice the employment rate follows an AR(1) process conditional on the labor demand shock z_{rt}. The coefficient on z_{rt} describes the initial effect of the shock: this is increasing in γ_1 but moderated by β_1; we call this parameter the shock amplification. And the persistence, identified by the coefficient on x_{rt-1}, is dampened by larger β_1 and β_2 as well as larger (more negative) γ_2; though the effect of γ_1 depends on the other parameters.

Our preferred estimates of the βs and γs[25] place the shock amplification (for the employment rate) at 0.62 and the AR(1) persistence at 0.55. As noted above, we do have concerns

[25]Our preferred estimates are the following: $\beta_1 = 0.664$ and $\beta_2 = 0.445$, based on the basic IV specification of column 4 of Table 2 (Panel A); and $\gamma_1 = 0.692$ and $\gamma_2 = -0.319$, based on the basic IV specification of column 4 of Table 7 (Panel A).

about our estimates of the employment equation - and γ_1 in particular. In Table 8, we compute the amplification and persistence for different combinations of γ_1 and γ_2, given our preferred estimates of β_1 and β_2. For comparison, we have included our own estimates of γ_1 and γ_2 in the table (0.692 and -0.319 respectively). The amplification is invariant to γ_2, though does vary markedly with γ_1. In contrast, the persistence is much more sensitive to γ_2 than γ_1. It is worth noting that turning the employment response off entirely (with both γ parameters taking zero) yields a very similar AR(1) parameter to our preferred estimates. This suggests that population growth, rather than employment growth, accounts for the bulk of the adjustment to local demand shocks: this is consistent with the findings of Blanchard and Katz (1992).

Consider a one-off permanent local demand shock: an innovation of 0.1 log points in z_{rt} for some area r at $t = 1$, from an initial position of a steady-state equilibrium. The impulse response is shown in Figure 3, based on our preferred estimates of the β and γ parameters. The large employment response (to population growth) greatly amplifies the impact on total stocks, with employment and population coming to rest at 0.18 log points above their original level. Of greater interest to this study is the impact on the employment rate. The population response moderates somewhat the initial impact of the innovation on the employment rate, though the effect clearly persists beyond one decade.

Having said that, the model alone clearly cannot match the magnitude of persistence in the data, as reported in Table 1, especially after the first lag or two. The impulse response function shows that equilibrium is largely restored in the model two decades after the shock. In contrast, in Table 1, the ACF of the local employment rate only reaches 0.5 by the sixth decadal lag.

5.3 Matching model to data

What explains this excess persistence in the data? In Section 3 above, we found little evidence that supply-side explanations play a substantial role. Instead, persistent demand-side factors appear more plausible. How much persistence in the unobserved demand shock z_{rt} is required to match the observed persistence in joblessness? Suppose that z_{rt} follows an AR(1) process:

$$z_{rt} = \lambda z_{rt-1} + u_{rt} \tag{23}$$

Substituting equation (23) for z_{rt} in (21) gives an AR(2) expression for x_{rt}:

$$(1 - \lambda L)(1 - \theta L) x_{rt} = \hat{u}_{rt} \tag{24}$$

where L is the lag operator, and

$$\theta = 1 - \frac{(1 - \gamma_1)\beta_2 - (1 - \beta_1)\gamma_2}{1 - \beta_1\gamma_1}$$

and

$$\hat{u}_{rt} = \left[\frac{1 - \beta_1}{1 - \beta_1\gamma_1}\right] u_{rt}$$

Notice in equation (24) that θ and λ are not separately identifiable using data on the employment rate x_{rt} alone. But, we can derive θ using our IV estimates of the β and γ parameters. And, we can then assign the remainder of the persistence to λ, the AR(1) parameter in the labor demand shock.

We plot the ACFs for different λ values in Figure 4, together with the ACF for the time-demeaned log employment rate, purged of observable amenity effects (the 10th row in Table 1). The $\lambda = 0$ line is the simulated ACF with a white noise shock. The model with $\lambda = 0$ accounts for about two thirds of the observed persistence at the first decadal lag, but its explanatory power declines substantially thereafter. To match the autocorrelation of the first lag, a λ of approximately 0.6 is required; and the third to sixth lags require a λ of 0.8. A higher order of persistence in z_{rt} would likely achieve a better fit.

Are these λ values realistic? In Figure 5, we plot the ACF of local employment growth.[26] There is clearly strong persistence, though this may conflate a range of supply and demand shocks (as well as feedback to these shocks). In addition, measurement error in the employment series may introduce significant biases.

We also plot the ACF for the time-demeaned Bartik shift-share, our instrument for changes in local employment. As a relatively exogenous measure of local demand growth, it may help guide our prior on the z_{rt} process. Figure 5 reveals strong persistence in this series also, similar to that of local employment growth. We estimate an AR(1) parameter for the Bartik shift-share of 0.68. This falls between the 0.6 and 0.8 bounds which, as Figure 4 shows, are required to match the employment rate ACF. This Bartik shock analogy yields a natural explanation for persistence in the demand shocks: a combination of long-run shifts in the industrial composition (as illustrated in Figure 6) and stickiness in local industrial

[26] Our employment growth series ends at the fifth lag, because our sample does not include local employment count data in 1940.

composition[27].

This story is consistent with our estimates of the (bias-corrected) ACF purged of CZ fixed effects, reported in rows 14-16 of Table 1. The fixed effects explain a relatively small fraction of the persistence at the first few decades (compared to the basic ACF in row 1), but potentially much more thereafter (depending on the assumption used to estimate the ACF). This makes sense if the fixed effects are largely serving to absorb the serial correlation in demand. As is clear from Figure 4, persistence in demand is more important for matching the ACF at larger lags: the model alone does a reasonable job at the initial lags.

The key idea comes from a well-known feature of ECMs: continuous shocks have a permanent effect on the deviation from equilibrium. To see this, suppose that z_{rt} in (19) is fixed at some constant z: every period, an area experiences the same increase or decrease in labor demand. In the steady-state, the local employment rate deviation x_{rt} must rest at an equilibrium x_r^*. Imposing $x_{rt} = x_r^*$ in equation (21) gives:

$$x_r^* = \frac{1 - \beta_1}{(1 - \gamma_1)\beta_2 - (1 - \beta_1)\gamma_2} z \tag{25}$$

so that areas facing continuous negative local demand shocks will have permanently lower employment rates. For our benchmark parameter values, the coefficient on z is 1.38. That is, a 1% permanent deviation in the local growth of labor demand yields a 1.38% deviation in the employment rate.

6 Other outcomes of interest

6.1 Wage and housing price effects

In this study, for the reasons explained above, we have focused on changes in quantities rather than prices. For completeness though, we document here the impact of local employment shocks on residualized wages and housing costs (the most important component of price differentials across areas). Specifically, we re-estimate the population equation (15), but simply change the dependent variable to log wage or price changes.

We construct residualized local wages and housing costs using the IPUMS census/ACS microdata samples from each cross-section since 1970. Although earnings and housing cost data are available before 1970, the census extract of 1960 does not include sub-state ge-

[27]The other side of this coin may be the positive demand shocks experienced by ideas-producing areas. See, for example, Moretti (2004), Glaeser (2005), Glaeser and Ponzetto (2007) and **?**.

ographical identifiers (see Appendix B). Our wage sample consists of full-time (at least 35 hours) and full-year (at least 40 weeks) employees aged 16-64, excluding those living in group quarters. We study the effect on weekly wages specifically, estimated by dividing annual labor earnings by weeks worked. Within each cross-section, we extract log wage residuals from a regression on detailed demographic characteristics.[28] We then estimate the average log residual within each CZ.[29]

We estimate residualized housing costs, separately for rents (for renters) and prices (for homeowners), in a similar way. We restrict the sample to houses and flats; we exclude farms, units with over 10 acres of land, and units with commercial use. Within each cross-section, we extract log housing cost residuals (separately for rents and prices) from a hedonic regression on a range of observed housing characteristics.[30] Then, as with wages, we estimate average log residuals within each CZ.

We report out estimates in Table 9. These are much more sensitive to specification than our population response estimates in Table 2. For example, the impact of the lagged employment rate on nominal wage growth varies not just in magnitude but also in sign, depending on whether we included CZ fixed effects. But generally, the estimates do suggest that nominal wages respond positively to the same factors as population growth. The same is true of rents and house prices.

Given that we do not have a model to interpret the estimates in Table 9, it is perhaps simpler to study the reduced form results. We report these in Table 10. The "long run" (two decade) effect of a Bartik shock on a variable can be derived by summing the coefficients on the current and lagged shocks. Notice the long run impact on all three variables is positive, though there is some negative feedback in the second decade - perhaps because a delayed labor supply response in the second decade puts downward pressure on wages. Notice also that the short run impact on prices is much larger than on rents, though the long run effects are similar. The short run overshooting of prices is presumably driven by expectations of

[28]Specifically, we control for age and age squared; four education effects, and interactions between these and the age quadratic; interactions between a gender dummy and all previously mentioned variables; and black, Hispanic and foreign-born dummies.

[29]As explained in Appendix B below, sub-state geographical identifiers vary by census cross-section. In general, these cannot precisely identify the CZ boundaries. Our strategy is to weight these log residuals using appropriate population allocations between these identifiers and the CZs; see Appendix B for further detail on these allocations.

[30]These consist of number of rooms (9 indicators); number of bedrooms (6 indicators); interaction between number of rooms and bedrooms; building age (up to 9 indicators, depending on cross-section); and indicators for kitchen, complete plumbing and condominium status. We also control for a house/flat dummy, as well as interactions between this and all previously mentioned variables.

asset price appreciation (see e.g. Gallin, 2008).

To assess the impact on real consumption wages, we need to know how much wages must grow relative to housing costs to keep individuals indifferent. One estimate might be the share of housing costs in total expenditure (about 0.25, according to Davis and Ortalo-Magne, 2011). By this reckoning, the "short run" (decadal) elasticity of real wages (deflated by housing rents) to a Bartik shock is 0.40[31], and the long run effect is 0.07. On the other hand, Albouy (2012) argues that a household requires a 0.65% increase in earnings to remain indifferent to a 1% increase in housing costs. His calculation takes into account housing's share of total expenditure (which he estimates at 21%), the effective federal tax rate on labor income adjusted for homeowner tax benefits (32%), the share of household income that depends on local wages (75%), and the fact that the cost-of-living differences across cities amount to a third of differences in local housing costs (net of homeowner tax benefits). By this reckoning, the short run effect on real wages would be approximately zero, and the long run effect would be negative. Clearly, the estimate of the real wage effect is very sensitive to this parametrization. This is one reason why we prefer to summarize local labor market conditions using the employment rate.

6.2 Commuting response to employment shocks

Until now, we have treated commuting zones as distinct local labor markets and assumed that workers who wish to work elsewhere must change residence. This seems reasonable, given that CZs are constructed to minimize the commuting flows between them. But, commuting flows between CZs are not negligible, and this raises the possibility that commuting itself may respond to economic shocks. We make use of cross-county commuting flow data published by the Census Bureau since 1970. Data has not been constructed for 2010, but is available for a pooled ACS sample of 2006-10 - so we use this instead.[32] The share of workers employed outside their CZ of residence grew from 3.4% in 1970 to 7.5% in 2006-10. As one would expect, there is considerable heterogeneity across CZs. In Figure 7, we plot the distribution across CZs of the share of workers commuting outside their CZ in 1970 and 2006-10. The

[31] This is equal to the wage effect minus one quarter of the rent effect.

[32] We take data for 1970 and 1980 from the BEA (*http://www.bea.gov/regional/histdata/releases/0405jtw/index.cfm*) and download flow data for 1980 and 1990 directly from the Census Bureau website (*http://www.census.gov/hhes/commuting/data/commuting.html*). Data for the ACS 2006-10 sample is available at *http://www.census.gov/population/metro/data/other.html*. Note that we estimate local commuting shares (that is, the share of workers employed outside their CZ of residence) based on data for *all* employed workers, rather than the 16-64 age group we usually study: age-disaggregated flows are unavailable in most cross-sections.

maximum share rose from 27% in 42% over the period.

Let n_{rt}^L be the log employment of area r residents who commute locally, with n_{rt} being the log employment of *all* area r residents. Our strategy is to re-estimate the population equation (20), but simply replace the dependent variable with $\Delta s_{rt}^L = \Delta n_{rt}^L - \Delta n_{rt}$; that is, the change in the log share of employed residents who work locally, i.e. do not commute across a CZ boundary.[33] By using the same right-hand side variables as in the population equation, we can directly compare the magnitude of the population and commuting responses to local employment conditions. The basic specification is the following:

$$\Delta s_{rt}^L = \delta_0 + \delta_1 \Delta n_{rt} + \delta_2 \left(n_{rt-1} - l_{rt-1}\right) + \delta_3' A_{rt} + \varepsilon_{rt} \tag{26}$$

where A_{rt} is a vector of controls identical to those we use above. Following large local employment growth, fewer workers in area r should be tempted to commute elsewhere (at significant cost): δ_1 should be positive. And a large initial employment rate should have a similar effect, so δ_2 should be positive.

An important concern is that commuting decisions are largely dependent on the employment situation in neighboring CZs, but these are themselves correlated with local conditions. Given this, we also estimate specifications of the form:

$$\Delta s_{rt}^L = \delta_0 + \delta_1 \Delta n_{rt} + \delta_2 \left(n_{rt-1} - l_{rt-1}\right) + \delta_3 \Delta n_{rt}^A + \delta_4 \left(n_{rt-1}^A - l_{rt-1}^A\right) + \delta_5' A_{rt} + \varepsilon_{ct} \tag{27}$$

where n_{rt}^A is the log of total employment across all adjacent CZs, and l_{rt}^A is the log of total population. 3 of our 722 CZs are islands[34], and we omit them for this exercise.

Clearly, employment growth and the initial employment rate in adjacent CZs are endogenous to Δs_{rt}^L. As a result, we require two more instruments to identify δ_3 and δ_4. For each CZ, we compute the (employment-weighted) average Bartik shift-share across adjacent CZs; and we use the current and lagged shift-shares as our instruments. We report our estimates in Table 11.

According to our basic IV estimates in column 1, the elasticity of the local commuting share with respect to local employment growth or the lagged employment rate is about 0.1. This is a substantial effect, given that the elasticity of population is around 0.5. These

[33] Given that most workers commute locally, variation in the log share is approximately equivalent to variation in the absolute share. But, the log specification yields a simpler interpretation of the estimates which we set out below.

[34] These are Martha's Vineyard (MA), Nantucket (MA) and San Juan Islands (WA).

effects are even larger, once we control (in column 2) for employment conditions in adjacent CZs. The coefficient on local employment growth reaches 0.18, and the coefficient on lagged employment rate 0.15. The effects of the adjacent CZs are negative and approximately equal in magnitude. The first stage results of the basic specification show considerable power in disentangling the effects of area r and the adjacent areas, with the appropriate instruments driving most of the variation in the corresponding endogenous variables.

The effect of the lagged employment rate is more than twice as large in the fixed effect specification (columns 3 and 4) and more than three times in the first differenced specification (columns 5 and 6), though the first stage results are less compelling: while the instruments retain much power, correlation patterns between the endogenous variables and instruments are more complex.

Broadly speaking, these results indicate that commuting plays an important role in the adjustment to local shocks, even when the areas are defined as commuting zones.

7 Conclusion

In this paper, we have explored the persistence in labor market opportunities across commuting zones in the US. We have argued for the use of the employment rate as a summary measure of local labor market opportunities, rather than the more conventional real wage. The two measures are equivalent as long as the labor supply or wage curve is not completely inelastic, which is plausible. And the quantity measure has two advantages. First, it suffers fewer measurement difficulties than real wages. And second, our estimates of the population response are directly informative of the speed of adjustment.

Joblessness across US commuting zones shows a high level of persistence even over a period of 60 years. We have claimed this cannot be explained by permanent (or highly persistent) differences in the characteristics of areas, whether due to demographic composition or amenities. Using a simple model, we derive a new ECM specification to assess local population changes, which explicitly accounts for the response to disequilibrium. We show that this model performs well empirically. The migration response is large, but not as large as commonly thought: adjustment to shocks is not complete within a decade. Over longer periods, we explain persistence in joblessness by serial correlation in the local demand shocks themselves. An adverse local demand shock is followed by population adjustment; but further negative shocks are likely to follow also, so the employment rate does not return to its original level. We have argued this persistence in local demand growth may be driven

by long-run changes in industrial composition, but the exact causes should be the subject of future research.

Appendix

A. Alternative local labor market models

The model presented in the main body of the paper is deliberately kept very simple to make clear the main ideas. But one might be concerned that some of the theoretical results are not robust to our assumptions. In this appendix, we sketch more elaborate models to address some of these concerns. We consider models with a non-traded goods sector that does employ labor, models with agglomeration effects, models with endogenous amenities, models based on frictional labor markets and models with heterogeneous skills.

First, it should be noted there is a simple explanation why the basic results are robust to these considerations. Notice that utility can be expressed as a function of the employment rate and the amenity, after combining the wage curve (6) and the utility equation (7). Essentially, for a given labor supply curve or "wage curve", the welfare of workers can be summarized by the position on the curve. This position can be expressed by either the real wage or employment rate, as long as there is some elasticity to the relationship. The validity of this argument is indepdent of how labour demand is modeled. And if amenities are endogenous, this will also be captured to the extent that they depend on variables that can be reduced to the employment rate. With this in mind, we next sketch more generalized versions of the simple model in the main text.

Non-traded goods sector that employs labor

Assume that the expenditure function in area r is given by:

$$y = E\left(P, P_r^n, P_r^h\right) A_r U_r \tag{28}$$

where P is the price of traded goods (treated as exogenous to the area), P_r^n is the local price of non-traded goods, P_r^h is the local price of housing, A_r is a local disamenity, and U_r is local utility. We assume that preferences are homothetic for simplicity, so that we do not need to keep track of income distribution within areas. Suppose that total income within an area r is given by:

$$Y_r = (1-\tau) W_r N_r + B (L_r - N_r) \tag{29}$$

so that income in a region derives from total labor income (which is taxed) and some benefits paid to residents who are not in work.

From (28) and (29), we can derive the following demands for the three different types of goods where lower-case denotes logarithms:

$$x_r = \frac{\partial e}{\partial p} + y_r \tag{30}$$

$$x_r^n = \frac{\partial e}{\partial p_r^n} + y_r \tag{31}$$

$$x_r^h = \frac{\partial e}{\partial p_r^h} + y_r \tag{32}$$

Now consider the production side of the economy. We assume that the quantity supplied of the three different goods is given by:

$$Q_{ir} = Q\left(P, P_r^n, P_r^h, W_r\right) Z_{ir} \tag{33}$$

where Z_{ir} is an area-specific industry shifter which is the source of the Bartik shock.

$$Q_r^n = Q^n\left(P, P_r^n, P_r^h, W_r\right) \tag{34}$$

$$Q_r^h = Q^s\left(P, P_r^n, P_r^h, W_r\right) \tag{35}$$

Associated with these supply functions are labor demand functions:

$$N_{ir} = N\left(P, P_r^n, P_r^h, W_r\right) Z_{ir}^n \tag{36}$$

$$N_r^n = N^n\left(P, P_r^n, P_r^h, W_r\right) \tag{37}$$

$$N_r^h = N^s\left(P, P_r^n, P_r^h, W_r\right) \tag{38}$$

Finally, to close the model, we need a labor supply curve or "wage curve". We simply assume that the employment rate is a function of the local real wage:

$$N_r = \sum_i N_{ir} + N_r^n + N_r^h = \left(W_r / E\left(P, P_r^n, P_r^h\right)\right)^\theta L_r \tag{39}$$

For a given local population, this allows us to solve for the locally-determined variables as a function of the exogenous variables. For the reasons given earlier, utility can be expressed as a function of the employment rate. The main insight from this model is that we would expect local employment to respond to local population directly and not just through any effect on the real wage. The reason is that non-traded goods demand has to be met from local employment. We use this insight to motivate our model for local employment growth.

Agglomeration effects in production

The most common way to model agglomeration effects is to assume that the scale of operations (perhaps measured by aggregate employment) affects productivity, so there is a spillover from the aggregate level of employment onto the demand for labor of individual firms. But, at the aggregate level, this produces a very similar labor demand curve as we have used (though if the agglomeration effects are very large, one does have to worry about stability of equilibria). Therefore, our basic equation relating utility to the employment rate will remain as before.

Endogenous amenities

It may be that the level of population or activity in an area affects the level of the amenities offered, e.g. by affecting the range of goods on offer, the crime rate, the level of social capital or population density itself (e.g. Glaeser et al., 2001; Glaeser and Redlick, 2009; Glaeser et al., 2009). The endogeneity of amenities may also amplify the impact of a given demand shock on welfare (see Diamond, 2013). But, if these endogenous responses can be summarized by the employment rate or real wage, this will still lead to equations similar to the ones we have used.

Labor markets with frictions

Beaudry et al. (2014b) use a local labor market model with frictions to investigate issues very similar to the ones we have considered. In this framework, the labor demand curve is replaced by a vacancy creation curve, and market tightness is measured using the ratio of unemployment to vacancies. But, there is a one-to-one relationship between the unemployment-vacancy ratio and the employment rate. So, the wage bargaining curve -

which gives a relationship between the real wage and the unemployment-vacancy ratio - can simply be translated to a wage curve like the one we have used.

On the labor demand side, the standard non-competitive model imposes constant returns to scale in production. And vacancies are created up to the point where the hiring cost, suitably amortized, is equal to the gap between the marginal product and the wage. The hiring cost is assumed to be a function of the unemployment-vacancy ratio which, as argued above, can be replaced by the employment rate; so this gives a negative relationship between the employment rate and the wage. This does differ from the labor demand curve we presented in the main body of the paper, as population appears directly in this relationship with an elasticity of 1 (see Beaudry et al., 2014a). But the elasticity will be less than 1 if one introduces some diminishing returns to labor on the production side (perhaps from imperfect competition); and we have shown above how, with non-traded goods, population will have a direct role in any case. So, models with frictions lead to very similar if not identical conclusions and specifications of empirical equations.

Heterogeneous labor

Much of the recent urban literature emphasizes differences by skill (e.g. Moretti, 2011; Notowidigdo, 2011; Diamond, 2013). If there are skill-specific labor supply curves, then one can derive equations for population growth in different skill groups as a function of skill-specific changes in employment and the lagged employment rate. But, the labor demand curve for one type of labor will depend on the wages of all types of labor, reflecting the complementarity or substituability between skill groups. This might alter the specification of the response of employment to population changes, though this is not our main focus of interest.

B. Manipulation of census and ACS data

Population and employment

Where possible, we take our population and employment data from the published county-level aggregates[35]. Published population counts (by age and gender) are based on 100% samples, while employment status counts (for employment, unemployed and inactive) are

[35]We extract these from the National Historical Geographic Information System (Minnesota Population Center, 2011). At the time of writing, county-level employment data for 1960 was unavailable on NHGIS. Instead, we take this data from the County Book of 1962, downloaded from ICPSR: *http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/12*.

usually based on samples of 15-20% (depending on the variable and year). The US Census Bureau did not implement a long form questionnaire in 2010, so we supplement data in that year with the American Community Survey (ACS) between 2009 and 2011; the ACS covers a 1% sample each year.

The published data do not report counts for all the demographic cells of interest.[36] Fortunately, micro-data samples from the US census and ACS are available for each cross-section. Our strategy is to disaggregate the published local population and employment counts (into components of interest) using CZ-specific shares estimated from the micro-data. For example, to impute the local population of college graduates aged 16-64 in a given year, we multiply the published local population count (in that age group) with the CZ-specific college graduate population share (estimated from micro-data).

Micro-data samples from the US census and ACS are made available for each cross-section by IPUMS (organized by **?**); and with the exception of 1960, these are accompanied by sub-state geographical identifiers. These identifiers vary across years[37], and it is not possible to perfectly identify commuting zones based on their boundaries. Following Autor and Dorn (2013) and Autor et al. (2013), we estimate population allocations of the geographical identifiers into CZs[38]; and we impute CZ outcomes by appropriately weighting outcomes for the available geographical identifiers with these allocations. Unfortunately, no sub-state identifiers are available in the 1960 micro-data. We impute CZ-level statistics of interest for 1960 by linear interpolation, using CZ-specific estimates for 1950 and 1970.

Industry shift-shares

To construct the industry shift-shares, we require local data on detailed industrial employment composition. We take this data from the IPUMS census extracts and ACS samples, restricting our sample to workers aged 16-64. IPUMS has created consistent 3-digit indus-

[36] In particular, county-level employment status counts are not disaggregated by age in the 1950, 1960 and 1970 censuses (except for the total labor force in 1970); these are simply reported for 14+ or 16+ groups. Also, disaggregations of both population and employment by education are not available in any year for the 16-64 age group.

[37] There are 467 State Economic Areas (SEAs) in the continental US in 1940 and 1950, 405 county groups in 1970, 1,148 county groups in 1980, 1,713 Public Use Microdata Areas (PUMAs) in 1990, 2,057 PUMAs in 2000 and the ACS until 2011, and 2,336 PUMAs in the ACS of 2012.

[38] We take county-SEA lookup tables from IPUMS (*https://usa.ipums.org/usa/resources/volii/sea_county_components.xls*) and use NHGIS county population counts to estimate CZ allocations. We also take our data for the 1970 and 1980 population allocations from IPUMS: see *https://usa.ipums.org/usa/resources/volii/1970cgcc.xls* and *https://usa.ipums.org/usa/resources/volii/cg98stat.xls* respectively. For the remaining years, we generate the allocations using the MABLE/Geocorr applications on the Missouri Census Data Center website: *http://mcdc.missouri.edu*.

try series based on both the 1950 and 1990 census scheme.[39] We use the consistent 1950 scheme to estimate industry shift-shares in the 1940s, 1950s, 1960s and 1970s; and we use the consistent 1990 scheme for all decades thereafter. As a precaution, since conversions of more disaggregated codes are inevitably less accurate, we aggregate all our industry data to the 2-digit level.

We use the same method outlined in the sub-section above to impute employment counts for each CZ (by industry), exploiting the available sub-state geographical identifiers. We impute employment counts in 1960 through the following steps. First, for each industry-year cell, we estimate the share of employment in each CZ in 1950 and 1970. Second, we impute values for 1960 by linear interpolation. Given the ad hoc nature of linear interpolation, these CZ shares will not sum to 1 (within industry-year cells); so we scale these shares by constant factors (within industry-year cells) to ensure that they do. And third, we multiple these shares with aggregate industry employment counts (estimated using the 1960 micro-data) to impute employment counts in CZ-industry cells.

Migrant shift-shares

We use the IPUMS micro-data to construct the migrant shift-share controls. This exercise requires a panel of CZ-origin-year cells. We include 79 origin countries in this panel. And we impute CZ-specific 1960 data in the same way outlined in the sub-section above.

In the main specifications, we restrict the inflow of new migrants $M^{new}_{o(-r)t}$ arriving between $t-1$ and t, in equation (17), to those aged 16-64 in year t. These quantities are simple to estimate between 1970 and 2010, because the census in these years and ACS include data on number of years in the US. For earlier census years, we impute $M^{new}_{o(-r)t}$ using cohort changes. For example, we estimate the inflow of new migrants of origin o in the 1950s as the difference between (1) the origin-specific stock of migrants in 1960 aged 16-64 and (2) the origin-specific stock of migrants in 1950 aged 6-54. The denominator L_{rt-1} in equation (17), the initial local population, is also restricted to individuals aged 16-64. But, we estimate the local share by migrant origin ϕ^{o}_{rt-1} using the entire sample (of all ages).

In the group-specific models, (by education, sex and age), we estimate the new migrant inflows $M^{new}_{o(-r)t}$ and initial population L_{rt-1} for the group of interest.

[39] See *https://usa.ipums.org/usa/volii/occ_ind.shtml*.

C. Unbiased estimator for the fixed effect ACF

In this appendix, we show how we derive an unbiased auto-correlation function for the time-demeaned log employment rate, x_{rt}, controlling for local fixed effects. Our data is limited to 7 time observations (over the period 1950-2010) and 722 areas, which we generalise in this exposition to T and R respectively. Suppose x_{rt} in region r is stationary with mean μ_r - we allow this to vary across regions. We are interested in modelling the average ACF across regions to let us assume, for simplicity, that C^n, the nth order covariance does not vary with region r i.e. we have:

$$C^n = E\left[\left(x_{rt} - \mu_r\right)\left(x_{rt-n} - \mu_r\right)\right] \tag{40}$$

for all r. But, C^n cannot be estimated dirrectly because μ_r is unknown.

Suppose we estimate μ_r using the sample mean:

$$\hat{\mu}_r = \frac{1}{T}\sum_{t=1}^{T} x_{rt} \tag{41}$$

We can use $\hat{\mu}_r$ to form a sample estimate of the covariance for region r :

$$\hat{C}_r^n = \frac{1}{T-n}\sum_{t=n+1}^{T}\left(x_{rt} - \hat{\mu}_r\right)\left(x_{rt-n} - \hat{\mu}_r\right) \tag{42}$$

for $n \leq T-1$. Since T is small, $\hat{C}_r^n$ is a biased estimator for C^n. But, we can derive the form of the bias. Specifically, taking expectations of (42):

$$\begin{aligned}
E\left(\hat{C}_r^n\right) &= E\left\{\frac{1}{T-n}\sum_{t=n+1}^{T}\left[\left(x_{rt} - \mu_r\right) - \left(\hat{\mu}_r - \mu_r\right)\right]\left[\left(x_{rt-n} - \mu_r\right) - \left(\hat{\mu}_r - \mu_r\right)\right]\right\} \\
&= E\left(\hat{\mu}_r - \mu_r\right)^2 + E\left\{\frac{1}{T-n}\sum_{t=n+1}^{T}\left(x_{rt} - \mu_r\right)\left(x_{rt-n} - \mu_r\right)\right\} \\
&\quad -E\left\{\frac{1}{T-n}\left(\hat{\mu}_r - \mu_r\right)\left[\sum_{t=n+1}^{T}\left(x_{rt} - \mu_r\right) + \sum_{t=n+1}^{T}\left(x_{rt-n} - \mu_r\right)\right]\right\} \\
&= C^n + E\left(\hat{\mu}_r - \mu_r\right)^2 \\
&\quad -2\frac{1}{T-n}E\left\{\left(\hat{\mu}_r - \mu_r\right)\sum_{t=n+1}^{T}\left(x_{rt} - \mu_r\right)\right\}
\end{aligned} \tag{43}$$

Notice that, from (41):

$$\hat{\mu}_r - \mu_r = \frac{1}{T}\sum_{t=1}^{T}(x_{rt} - \mu_r) \tag{44}$$

so that:

$$\begin{aligned} E(\hat{\mu}_r - \mu_r)^2 &= \frac{1}{T^2}E\left[\sum_{t=1}^{T}(x_{rt} - \mu_r)\right]^2 \\ &= \frac{1}{T^2}E\left[\sum_{t=1}^{T}(x_{rt} - \mu_r)^2 + 2\sum_{s=1}^{T-1}\sum_{t=s+1}^{T}(x_{rt} - \mu_r)(x_{rt-s} - \mu_r)\right] \\ &= \frac{1}{T}C^0 + \frac{2}{T^2}\sum_{s=1}^{T-1}(T-s)C^s \end{aligned} \tag{45}$$

is a linear function of the covariances.

Also, notice that:

$$\begin{aligned} 2\frac{1}{T-n}E\left\{(\hat{\mu}_r - \mu_r)\sum_{t=n+1}^{T}(x_{rt} - \mu_r)\right\} &= 2\frac{1}{T}\frac{1}{T-n}E\left\{\sum_{s=1}^{T}(x_{rs} - \mu_r)\sum_{t=1}^{T-n}(x_{rt} - \mu_r)\right\} \\ &= 2\frac{1}{T}\frac{1}{T-n}E\left\{\left[\sum_{t=1}^{T-n}(x_{rt} - \mu_r)\right]^2 + \Omega_r^n\right\} \\ &= 2\frac{1}{T}\frac{1}{T-n}E\left\{\left[\sum_{t=1}^{T}(x_{rt} - \mu_r)\right]^2 - \Omega_r^n\right\} \end{aligned} \tag{46}$$

where Ω^n is an unknown quantity, which can be derived by equating the final two lines:

$$\begin{aligned} E(\Omega_r^n) &= \frac{1}{2}E\left\{\left[\sum_{t=1}^{T}(x_{rt} - \mu_r)\right]^2 - \left[\sum_{t=1}^{T-n}(x_{rt} - \mu_r)\right]^2\right\} \\ &= \frac{1}{2}\left[TC^0 + 2\sum_{s=1}^{T-1}(T-s)C^s\right] - \frac{1}{2}\left[(T-n)C^0 + 2\sum_{s=1}^{T-n-1}(T-n-s)C^s\right] \\ &= \frac{n}{2}C^0 + \sum_{s=1}^{T-1}(T-s)C^s - \sum_{s=1}^{T-n-1}(T-n-s)C^s \end{aligned} \tag{47}$$

and substituting this back into (46):

$$2\frac{1}{T-n}E\left\{(\hat{\mu}_r-\mu_r)\sum_{t=n+1}^{T}(x_{rt}-\mu_r)\right\} = 2\frac{1}{T}\frac{1}{T-n}E\left[\sum_{t=1}^{T}(x_{rt}-\mu_r)\right]^2 \quad (48)$$
$$-\frac{n}{2}C^0-\sum_{s=1}^{T-1}(T-s)\,C^s+\sum_{s=1}^{T-n-1}(T-n-s)\,C^s$$

Next, susbtituting (45) and (48) into (43) gives:

$$E\left(\hat{C}_r^n\right) = C^n+E\left(\hat{\mu}_r-\mu_r\right)^2-2\frac{T}{T-n}E\left(\hat{\mu}_r-\mu_r\right)^2 \quad (49)$$
$$+\frac{1}{T}\frac{1}{T-n}\left[nC^0+2\sum_{s=1}^{T-1}(T-s)\,C^s-2\sum_{s=1}^{T-n-1}(T-n-s)\,C^s\right] \quad (50)$$
$$= C^n-\frac{1}{T}\frac{T+n}{T-n}\left[C^0+\frac{2}{T}\sum_{s=1}^{T-1}(T-s)\,C^s\right] \quad (51)$$
$$+\frac{1}{T}\frac{1}{T-n}\left[nC^0+2\sum_{s=1}^{T-1}(T-s)\,C^s-2\sum_{s=1}^{T-n-1}(T-n-s)\,C^s\right] \quad (52)$$
$$= C^n-\frac{1}{T}\frac{1}{T-n}\left[TC^0+2\frac{2T+n}{T}\sum_{s=1}^{T-1}(T-s)\,C^s-2\sum_{s=1}^{T-n-1}(T-n-s)\,C^s\right] \quad (53)$$

This shows that the expectation of the biased covariance estimators $\hat{C}_r^n$ are linear functions of the true covariances C^n.[40] That is, there exists a $T\times T$ square matrix A such that:

$$E\left(\hat{C}_r\right)=AC \quad (54)$$

where $\hat{C}_r$ is a T-length vector of the sample covariances for region r, $\hat{C}_r^n$; and similarly, C is a vector of the true covariances C^n. In the context of (54), a natural way to derive an unbiased estimator for C would be to invert the matrix A. The problem is that A does not have full rank. It is easiest to get the intuition for this if $T=2$. In this case one cannot separately identify the variance and the first-order covariance as, effectively the only useful informatio is in $x_{r2}-x_{r1}$. Similarly, T observations are insufficient to identify $T-1$ variance/covariance parameters. One further restriction on the covariances is required for identification. We assume that $C^{T-1}=\pi C$, with $\pi<1$. This implies:

[40] In particular, notice that $C^0=\frac{T}{T-1}E\left(\hat{C}_r^0\right)$ if observations are independent, from which the standard formula for deriving an unbiased estimate of the variance follows.

$$\begin{pmatrix} E\left[\hat{C}_r^0\right] \\ \vdots \\ E\left[\hat{C}_r^{T-2}\right] \end{pmatrix}_{(T-1)\times 1} = \begin{pmatrix} A_1 & a_1 \end{pmatrix}_{(T-1)\times T} \begin{pmatrix} C^0 \\ \vdots \\ C^{T-2} \\ \pi C^{T-2} \end{pmatrix}_{T\times 1} \tag{55}$$

where

$$A_1 = \begin{pmatrix} A[0,0] & \cdots & A[0,T-2] \\ \vdots & \ddots & \vdots \\ A[T-2,0] & \cdots & A[T-2,T-2] \end{pmatrix}_{(T-1)\times(T-1)} \tag{56}$$

is the top left submatrix of A, exlcuding the final column and final row. And

$$a_1 = \begin{pmatrix} A[0,T-1] \\ \vdots \\ A[T-2,T-1] \end{pmatrix}_{(T-1)\times 1} \tag{57}$$

is the final column of A, excluding the final row. And so:

$$\begin{pmatrix} E\left[\hat{C}_r^0\right] \\ \vdots \\ E\left[\hat{C}_r^{T-2}\right] \end{pmatrix}_{(T-1)\times 1} = \begin{pmatrix} A_1 & \pi a_1 \end{pmatrix}_{(T-1)\times T} \begin{pmatrix} C^0 \\ \vdots \\ C^{T-2} \\ C^{T-2} \end{pmatrix}_{T\times 1} \tag{58}$$

which implies:

$$\begin{pmatrix} E\left[\hat{C}_r^0\right] \\ \vdots \\ E\left[\hat{C}_r^{T-2}\right] \end{pmatrix}_{(T-1)\times 1} = \left[A_1 + \begin{pmatrix} 0 & \cdots & 0 & \\ \vdots & & \vdots & \pi a_1 \\ 0 & \cdots & 0 & \end{pmatrix} \right]_{(T-1)\times(T-1)} \begin{pmatrix} C^0 \\ \vdots \\ C^{T-2} \end{pmatrix}_{(T-1)\times 1} \tag{59}$$

where the square matrix in (59) is the sum of (1) A_1 and (2) a $(T-1)\times(T-1)$ square matrix with πa_1 in the final column and 0s in the remaining columns. This yields a set of unbiased estimator $\tilde{C}^n$ for the true covariances of the form:

$$\begin{pmatrix} \tilde{C}^0 \\ \vdots \\ \tilde{C}^{T-2} \end{pmatrix} = \left[A_1 + \begin{pmatrix} 0 & \cdots & 0 & \\ \vdots & & \vdots & \pi a_1 \\ 0 & \cdots & 0 & \end{pmatrix} \right]^{-1}_{(T-1)\times(T-1)} \sum_{r=1}^{R} \begin{pmatrix} \hat{C}_r^0 \\ \vdots \\ \hat{C}_r^{T-2} \end{pmatrix} \tag{60}$$

with the highest order covariance estimator $\tilde{C}^{T-1}$ set to $\pi\tilde{C}^{T-2}$. We sum across regions to take the average of the covariances. It follows that $\tilde{C}$ is an unbiased estimate of C and consistent as one takes R to infinity.

The average nth order ACF is then estimated as:

$$ACF^n = \frac{\tilde{C}^n}{\tilde{C}^0} \tag{61}$$

For large R and small T, this is a consistent estimate of the true ACF.

D. Robustness checks for population response

In Table 12, we study the robustness of the IV population responses estimates from equation (15). Panel A reports a range of specifications with different controls, weighting the regressions by the lagged local population share (as we do in the estimates above). Notice that columns 7, 9 and 10 in Panel A are identical to columns 4, 5 and 6 respectively of Table 2 above. Panel B reports unweighted estimates of the same specifications.

The coefficient on contemporaneous employment growth varies little with the choice of controls or weighting: our estimates range from 0.604 to 0.819, with the standard error hovering around 0.03. In contrast, the coefficient on the lagged employment ratio varies somewhat more, especially among the weighted estimates of Panel A. In particular, the estimate in column 1 with no controls is 0.268, compared to 0.445 in our preferred specification (with no fixed effects or first differencing). Most of the difference is driven by the inclusion of the climate controls of column 3. The interactions between the time-invariant amenity controls and year effects in column 7 have little effect. The omission of the lagged ECM term (and its lagged Bartik instrument) in column 8 yields a larger coefficient on the change in employment: 0.819 compared to 0.664. This is a consequence of the serial correlation in the Bartik instrument. Finally, as we note in the main text, the fixed effects and first differenced estimates in columns 9 and 10 are significantly larger than the basic estimates in column 7.

In Panel B, the coefficient on the lagged employment ratio generally becomes smaller as controls are included. The estimate with no controls in column 1 is 0.580, and the full set of controls in column 7 yields an estimate of 0.383. Notice this is not much different to

our weighted estimate (0.448) in Panel A. This suggests the responsiveness of population is not markedly different in larger cities. Interestingly also, including fixed effects or first differencing (columns 9 and 10) makes much less difference in the unweighted specifications, with those estimates falling at 0.572 and 0.691 respectively.

In Figure 8, we present the estimates from column 7 (our preferred specification) graphically. This exercise is useful in demonstrating that the estimates are not driven by outliers. These plots follow the logic of the Frisch-Waugh theorem, but applied to 2SLS. In particular, we use the first stage regressions to predict a full set of values for each endogenous variable. Then, on the y-axis of the first panel, we plot the residuals from a regression of population growth on the predicted lagged employment growth, together with all exogenous variables. On the x-axis, we plot the residuals of a regression of employment growth on the same set of explanatory variables. In the second panel, we repeat the exercise for the lagged employment rate. Notice the standard errors of the best-fit slopes do not correspond to those in Table 12; this is because this naive estimator does not account for the standard error in the first stage. In any case, it is clear that the result is not driven by outliers.

E. Supplementary first stage estimates

Table 13 reports the first stage regressions which accompany the IV estimates in Tables 5 and 6. These are estimates of the population response equation (15), with the endogenous population and employment variables restricted to particular education, gender and age groups.

Tables and figures

Table 1: The Auto-Correlation Function for the employment rate

	Employment rate variant	Lag					
		1	2	3	4	5	6
(1)	Raw Data (time-demeaned)	0.84	0.78	0.71	0.60	0.55	0.52
	Sub-samples						
(2)	Years 1950-80	0.82	0.78	0.72	-	-	-
(3)	Years 1980-10	0.85	0.73	0.73	-	-	-
(4)	Labor force	0.53	0.45	0.46	0.38	0.36	0.28
(5)	College graduate	0.54	0.22	0.17	0.11	0.06	0.05
(6)	Non-graduate	0.82	0.74	0.66	0.53	0.46	0.44
(7)	Male	0.74	0.68	0.65	0.53	0.45	0.25
(8)	Female	0.90	0.78	0.68	0.55	0.41	0.42
(9)	Composition-adjusted	0.84	0.76	0.68	0.65	0.51	0.50
(10)	CZ amenity controls	0.84	0.78	0.69	0.57	0.49	0.41
(11)	Within-state	0.76	0.66	0.53	0.36	0.29	0.19
(12)	Collapsed to state	0.82	0.75	0.69	0.57	0.53	0.51
	Within-CZ						
(13)	Unadjusted	0.25	-0.02	-0.26	-0.53	-0.42	-0.49
(14)	Bias-corrected: $\pi = 0.9$	0.81	0.72	0.65	0.53	0.49	0.44
(15)	Bias-corrected: $\pi = 0.5$	0.69	0.55	0.43	0.22	0.16	0.08
(16)	Bias-corrected: $\pi = 0$	0.67	0.51	0.38	0.15	0.09	0

This table summarizes autocorrelation functions of the time-demeaned log employment rate, across five decadal lags. In general, these are estimated as the ratio of the lag-specific autocovariance to the product of the current and lagged standard deviations (weighted by CZ population share), across all CZs. Rows 2-8 estimate ACFs for particular subsamples of the data. Row 9 reports the ACF for the log employment rate adjusted for local demographic composition; the construction of this variable is described in footnote 19. As explained in the footnote, we are unable to construct this data for 1960, so we omit that year from our data series. Consequently, the sample used in the estimation of the ACF in row 9 is somewhat different to the other rows, though we take comfort from the stability in our estimates over time (as indicated by rows 2 and 3). Row 10 purges the log employment rate of variation in observed amenities, as described in footnote 20. In row 11, we purge the log employment rate of state fixed effects (as well as time effects). Some CZs do straddle state boundaries, so we allocate these to the state accounting for the largest population share of the CZ. In row 12, we estimate ACFs for the (time-demeaned) log employment rate of the 48 states of the Continental US (rather than the 722 CZs). Row 13 purges the log employment rate of CZ fixed effects (in the same way that row 11 removed state fixed effects). But, given the short panel, these estimates are biased. We correct for this bias following the procedure described in Appendix C. This procedure requires one identifying assumption. We fixed the ratio π of the sixth to fifth autocorrelation, and report estimates for different π in rows 14-16.

Table 2: Basic estimates for population response

PANEL A: OLS and IV

	OLS			IV		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Δ log emp 16-64	0.793***	0.790***	0.822***	0.664***	0.813***	0.715***
	(0.013)	(0.014)	(0.012)	(0.026)	(0.037)	(0.030)
Lagged log emp rate 16-64	0.198***	0.576***	1.047***	0.445***	1.037***	0.796***
	(0.015)	(0.033)	(0.030)	(0.051)	(0.156)	(0.101)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

PANEL B: First stage

	Δ log emp 16-64			Lagged log emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Current Bartik shock	1.147***	1.105***	1.008***	0.011	-0.146***	-0.051
	(0.069)	(0.074)	(0.080)	(0.040)	(0.033)	(0.032)
Lagged Bartik shock	-0.043	-0.157**	-0.223**	0.534***	0.217***	0.244***
	(0.060)	(0.064)	(0.087)	(0.052)	(0.034)	(0.025)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

This table reports OLS and IV estimates of β_1 and β_2 in the population response equation (15), as well as the first stage estimates (in Panel B), across 722 CZs and six (decadal) time periods. In the basic specification (columns 1 and 4), the dependent variable is the log decadal change in the CZ population, and the endogenous variables are the log change in local employment and the lagged log employment-population ratio, with all stocks corresponding to individuals aged 16-64. These are instrumented by the current and lagged Bartik industry shift-shares respectively, as defined in equation (16). We also control for a full set of time effects, three climate variables (the maximum January and July termperatures, and mean July relative humidity), a dummy for the presence of coastline, the log population density in 1900, the log distance to the closest CZ centroid, and the immigrant shift-share (as described in equation (17)); all these controls, excluding the immigrant shift-share, are also interacted with all the time effects. We report estimates from a fixed effects specification in columns 2 and 5, where we control for a full set of CZ dummies. And columns 3 and 6 report a first differenced specification, where all variables in the basic specification (including the instruments) are differenced; so the dependent variable in the double difference in log population (we lose one decade in our sample through this transformation). Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1. ****.

Table 3: Basic estimates for labor force response

PANEL A: OLS and IV

	OLS			IV		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Δ log emp 16-64	0.947***	0.954***	0.955***	0.883***	0.980***	0.886***
	(0.006)	(0.007)	(0.006)	(0.014)	(0.040)	(0.020)
Lagged log emp rate 16-64	0.453***	0.947***	1.330***	1.253***	2.660***	1.061***
	(0.022)	(0.047)	(0.033)	(0.257)	(0.890)	(0.174)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

PANEL B: First stage

	Δ log emp 16-64			Lagged log emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Current Bartik shock	1.147***	1.105***	1.008***	-0.004	-0.040***	0.007
	(0.069)	(0.074)	(0.080)	(0.010)	(0.010)	(0.017)
Lagged Bartik shock	-0.043	-0.157**	-0.223**	0.075***	0.037**	0.076***
	(0.060)	(0.064)	(0.087)	(0.011)	(0.016)	(0.014)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

This table reports OLS and IV estimates of the labor force response to employment shocks, as well as the first stage estimates (in Panel B), across 722 CZs and six (decadal) time periods. The empirical specifications are identical to those in Table 2, based on equation (15), except the endogenous stock variables are restricted to labor force participants. That is, the dependent variable is the log change in the local active labor force, and the endogenous variables are the log employment change and the lagged log employment rate (measured here as the ratio of employed to all participants), with all stocks corresponding to individuals aged 16-64. See the notes of Table 2 for further details. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Population estimates before and after 1980

	1950-1980			1980-2010		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
OLS						
Δ log emp 16-64	0.805***	0.839***	0.867***	0.772***	0.697***	0.758***
	(0.013)	(0.020)	(0.015)	(0.022)	(0.024)	(0.022)
Lagged log emp rate 16-64	0.195***	0.911***	1.156***	0.204***	0.717***	0.848***
	(0.022)	(0.049)	(0.041)	(0.018)	(0.060)	(0.052)
IV						
Δ log emp 16-64	0.734***	0.871***	0.849***	0.524***	0.545***	0.450***
	(0.034)	(0.040)	(0.041)	(0.032)	(0.047)	(0.073)
Lagged log emp rate 16-64	0.378***	0.956***	1.027***	0.482***	0.657***	0.312**
	(0.068)	(0.150)	(0.120)	(0.078)	(0.111)	(0.149)
FS: Δ *log emp 16-64*						
Current Bartik shock	1.133***	1.055***	1.023***	1.210***	0.908***	0.708***
	(0.091)	(0.130)	(0.104)	(0.092)	(0.193)	(0.178)
Lagged Bartik shock	-0.051	-0.094	-0.056	0.024	-0.566**	-0.558***
	(0.079)	(0.137)	(0.093)	(0.127)	(0.225)	(0.201)
FS: *Lagged log emp rate 16-64*						
Current Bartik shock	0.005	-0.126***	-0.106***	0.074	0.146	0.184***
	(0.035)	(0.039)	(0.031)	(0.087)	(0.096)	(0.071)
Lagged Bartik shock	0.476***	0.239***	0.226***	0.789***	0.523***	0.462***
	(0.063)	(0.057)	(0.043)	(0.089)	(0.082)	(0.056)
Observations	2,166	2,166	1,444	2,166	2,166	1,444

Like Table 2, this table reports OLS, IV and first stage estimates of the population response equation (15) across 722 CZs, though this time separately for the period 1950-1980 (in the first three columns) and 1980-2010 (in the final three). See the notes of Table 2 for further details. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Estimates by skill

	College graduates			Non-graduates		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
OLS						
Δ log emp 16-64	0.943***	0.947***	0.951***	0.770***	0.764***	0.796***
	(0.007)	(0.007)	(0.005)	(0.015)	(0.017)	(0.014)
Lagged log emp rate 16-64	0.702***	0.781***	0.912***	0.193***	0.571***	1.018***
	(0.062)	(0.053)	(0.016)	(0.016)	(0.033)	(0.032)
IV						
Δ log emp 16-64	0.865***	0.829***	0.838***	0.636***	0.805***	0.684***
	(0.046)	(0.083)	(0.055)	(0.028)	(0.046)	(0.032)
Lagged log emp rate 16-64	0.729***	0.610*	0.628**	0.452***	1.113***	0.806***
	(0.155)	(0.346)	(0.261)	(0.056)	(0.197)	(0.114)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

Like Table 2, this table reports OLS and IV estimates of the population response equation (15) across 722 CZs and six (decadal) time periods, though this time separately for college graduates (with at least four years of college) and non-graduates. The endogenous variables in these specifications are education-specific. For example, for college graduates, we regress the log decadal change in the CZ graduate population on the log change in graduate employment and the lagged log graduate employment rate. Similarly, the immigrant shift-share is constructed to predict the contribution of new migrants to the education-specific local population (see Appendix B). Otherwise, these specifications are identical to those in Table 2; see the notes below that table for further details. In particular, the instruments are the same as in Table 2: the current and lagged Bartik shift-share in each case. First stage estimates are reported in the Appendix E. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 6: IV estimates by sex and age

Age group	16-64		16-24	25-44	45-64
	Male	Female	All	All	All
	(1)	(2)	(3)	(4)	(5)
Basic specification					
Δ log emp	0.732***	0.531***	0.616***	0.751***	0.478***
	(0.020)	(0.066)	(0.024)	(0.028)	(0.041)
Lagged log emp rate	0.430***	0.445***	0.414***	0.530***	0.603***
	(0.058)	(0.051)	(0.035)	(0.067)	(0.086)
CZ fixed effects					
Δ log emp	0.935***	0.430***	0.761***	0.851***	0.753***
	(0.072)	(0.075)	(0.068)	(0.026)	(0.055)
Lagged log emp rate	1.083***	0.871***	1.014***	1.241***	1.031***
	(0.230)	(0.137)	(0.237)	(0.172)	(0.125)
Observations	4,332	4,332	4,332	4,332	4,332
First differences					
Δ log emp	0.736***	0.537***	0.586***	0.805***	0.603***
	(0.039)	(0.040)	(0.037)	(0.026)	(0.041)
Lagged log emp rate	0.642***	0.830***	0.522***	1.073***	0.871***
	(0.136)	(0.075)	(0.129)	(0.107)	(0.072)
Observations	3,610	3,610	3,610	3,610	3,610

Like Table 2, this table reports OLS and IV estimates of the population response equation (15) across 722 CZs and six (decadal) time periods, though this time separately for gender and age groups (16-24, 25-44 and 45-64). The endogenous variables in these regressions are group-specific. For example, in column 1, we regress the log decadal change in the CZ male population on the log change in male employment and the lagged log male employment rate. Similarly, the immigrant shift-share is constructed to predict the contribution of new migrants to the group-specific population (see Appendix B). Otherwise, these specifications are identical to those in Table 2; see the notes below that table for further details. In particular, the instruments are the same as in Table 2: the current and lagged Bartik shift-share in each case. First stage estimates are reported in the Appendix E. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Estimates for employment response

PANEL A: OLS and IV

	OLS			IV		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Δ log pop 16-64	1.012***	0.985***	1.011***	0.692***	-7.847	2.076***
	(0.011)	(0.018)	(0.016)	(0.075)	(28.625)	(0.688)
Lagged log emp rate 16-64	-0.148***	-0.727***	-1.264***	-0.319***	3.386	-1.316***
	(0.014)	(0.044)	(0.035)	(0.068)	(13.677)	(0.358)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

PANEL B: First stage

	Δ log pop 16-64			Lagged log emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Basic specification						
Max temp January	0.385***			0.022		
	(0.086)			(0.059)		
Max temp January * time	-0.014	-0.017	-0.008***	-0.048***	-0.043***	-0.009***
	(0.015)	(0.016)	(0.003)	(0.014)	(0.012)	(0.003)
Lagged Bartik shock	0.216***	0.113**	0.044	0.529***	0.218***	0.246***
	(0.056)	(0.053)	(0.062)	(0.049)	(0.034)	(0.024)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

This table reports OLS and IV estimates of γ_1 and γ_2 in the employment response equation (18), as well as the first stage estimates (in Panel B), across 722 CZs and six (decadal) time periods. In the basic specification (columns 1 and 4), the dependent variable is the log decadal change in CZ employment, and the endogenous variables are the log change in local population and the lagged log employment-population ratio, with all stocks corresponding to individuals aged 16-64. As in the population equation (see Table 2), we instrument the lagged employment rate with the lagged Bartik shift-share. And we instrument local population growth with two variables: the (time-invariant) maximum January tenmperature and its interaction with a time trend. We also control for a full set of time effects, two climate variables (the maximum July termperature and mean July relative humidity), a dummy for the presence of coastline, the log population density in 1900, the log distance to the closest CZ centroid, the immigrant shift-share (as described in equation (17)) and the current Bartik shift-share; all these controls, excluding the immigrant and Bartik shift-shares, are also interacted with all the time effects. We report estimates from a fixed effects specification in columns 2 and 5, where we control for a full set of CZ dummies. And columns 3 and 6 report a first differenced specification, where all variables in the basic specification (including the instruments) are differenced; so the dependent variable in the double difference in log population (we lose one decade in our sample through this transformation). Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Sensitivity of computed shock amplification and persistence parameters

γ_1 value	γ_2 value			
	0	-0.319	-0.6	-0.9
	Shock amplification			
0	0.336	0.336	0.336	0.336
0.3	0.420	0.420	0.420	0.420
0.692	0.622	0.622	0.622	0.622
0.9	0.835	0.835	0.835	0.835
	AR(1) persistence			
0	0.555	0.448	0.353	0.252
0.3	0.611	0.477	0.359	0.233
0.692	0.747	0.548	0.373	0.187
0.9	0.890	0.623	0.388	0.138

This table computes the parameters, in equation (21), on (1) the demand shock z_{rt} (the shock amplification) and (2) the lagged log employment rate deviation (the AR(1) persistence), for different values of γ_1 and γ_2. In each instance, we set β_1 and β_2 at 0.664 and 0.445 respectively, our preferred estimates from the basic specification of the population equation (15) in Table 2. Notice the $\gamma_1 = 0.692$ and $\gamma_2 = -0.319$ cases correspond to our preferred estimates from the basic specification of the employment equation (18) in Table 7.

Table 9: Wage and price effects: OLS and IV

	OLS			IV		
	Δ log wage	Δ log rent	Δ log price	Δ log wage	Δ log rent	Δ log price
	(1)	(2)	(3)	(4)	(5)	(6)
Basic specification						
Δ log emp 16-64	0.227***	0.365***	0.822***	0.575***	0.892***	1.497***
	(0.023)	(0.025)	(0.096)	(0.068)	(0.081)	(0.181)
Lagged log emp rate 16-64	-0.038*	0.020	-0.134*	-0.533***	-0.230	-0.909***
	(0.021)	(0.031)	(0.076)	(0.124)	(0.147)	(0.277)
CZ fixed effects						
Δ log emp 16-64	0.540***	0.729***	1.657***	0.906***	1.195***	1.806***
	(0.043)	(0.053)	(0.158)	(0.103)	(0.096)	(0.258)
Lagged log emp rate 16-64	0.465***	0.471***	0.471*	0.659*	0.904**	-0.042
	(0.073)	(0.137)	(0.285)	(0.377)	(0.381)	(0.640)
Observations	2,888	2,888	2,888	2,888	2,888	2,888
First differences						
Δ log emp 16-64	0.686***	0.910***	2.062***	1.107***	1.277***	2.203***
	(0.047)	(0.052)	(0.176)	(0.134)	(0.106)	(0.353)
Lagged log emp rate 16-64	0.622***	0.746***	0.929**	1.471***	1.003**	0.709
	(0.127)	(0.174)	(0.433)	(0.367)	(0.470)	(1.167)
Observations	2,166	2,166	2,166	2,166	2,166	2,166

This table reports OLS and IV estimates of the response of wages and housing costs to employment shocks, across 722 CZs and four (decadal) time periods beginning in 1970. We cannot construct this data in 1960, since the census microdata of that year does not include sub-state geographical identifiers (see Appendix B). The empirical specifications are identical to those in Table 2, based on the population response equation (15), except we change the dependent variable to the local change in the log residualized wage (in columns 1 and 4), residualized housing rent (columns 2 and 5) and residualized house price (3 and 6). The residualization method for each variable is described in Section 6.1. The two instruments are the current and lagged Bartik shift-share. See the notes under Table 2 for further details on the empirical specification. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 10: Wage and price effects: reduced form

	Δ log wage (1)	Δ log rent (2)	Δ log price (3)
Basic specification			
Current Bartik shock	0.647***	1.001***	1.683***
	(0.095)	(0.103)	(0.204)
Lagged Bartik shock	-0.362***	-0.15	-0.611***
	(0.113)	(0.092)	(0.161)
CZ fixed effects			
Current Bartik shock	0.926***	1.223***	1.805***
	(0.122)	(0.163)	(0.373)
Lagged Bartik shock	-0.086	-0.104	-0.561*
	(0.212)	(0.180)	(0.305)
Observations	2,888	2,888	2,888
First differences			
Current Bartik shock	1.083***	1.214***	2.042***
	(0.140)	(0.165)	(0.447)
Lagged Bartik shock	-0.036	-0.246	-0.725
	(0.239)	(0.243)	(0.586)
Observations	2,166	2,166	2,166

This table reports reduced form estimates of the response of wages and housing costs to employment shocks, which correspond to the OLS and IV estimates in Table 9. These are based on the population response equation (15), estimated in Table 2, except we change the dependent variable to the local change in the log residualized wage (in columns 1 and 4), residualized housing rent (columns 2 and 5) and residualized house price (3 and 6). The residualization method for each variable is described in Section **??**. The two instruments are the current and lagged Bartik shift-share. The sample covers 722 CZs and four (decadal) time periods beginning in 1970. We cannot construct this data in 1960, since the census microdata of that year does not include sub-state geographical identifiers (see Appendix B). See the notes under Table 2 notes for details on the included controls. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 11: Estimates for changes in log share of employed residents working locally

	Basic		FE		FD	
	(1)	(2)	(3)	(4)	(5)	(6)
IV						
Δ log emp 16-64	0.114***	0.182***	0.096***	0.196***	0.130***	0.210***
	(0.024)	(0.041)	(0.023)	(0.040)	(0.031)	(0.054)
Lagged log emp rate 16-64	0.107***	0.154***	0.218***	0.493***	0.371***	0.957***
	(0.024)	(0.042)	(0.073)	(0.113)	(0.118)	(0.248)
Δ log emp 16-64 (adj)		-0.154***		-0.187***		-0.171***
		(0.038)		(0.044)		(0.047)
Lagged log emp rate 16-64 (adj)		-0.157**		-0.382***		-0.784***
		(0.071)		(0.111)		(0.246)
FS: Δ log emp 16-64						
Current Bartik shock	1.121***	1.025***	1.000***	0.785***	0.911***	0.782***
	(0.101)	(0.127)	(0.132)	(0.160)	(0.143)	(0.162)
Lagged Bartik shock	0.01	0.112	-0.304**	-0.223	-0.425**	-0.296
	(0.094)	(0.102)	(0.143)	(0.158)	(0.186)	(0.192)
Current Bartik shock (adj)		0.201		0.617***		0.382**
		(0.174)		(0.163)		(0.160)
Lagged Bartik shock (adj)		-0.282**		-0.142		-0.279**
		(0.122)		(0.138)		(0.135)
FS: Lagged emp rate 16-64						
Current Bartik shock	-0.007	0.017	0.031	0.053	0.051	0.032
	(0.119)	(0.125)	(0.064)	(0.071)	(0.048)	(0.050)
Lagged Bartik shock	0.689***	0.681***	0.287***	0.212***	0.295***	0.212***
	(0.089)	(0.098)	(0.039)	(0.043)	(0.033)	(0.038)
Current Bartik shock (adj)		-0.066		-0.056		0.022
		(0.115)		(0.072)		(0.054)
Lagged Bartik shock (adj)		0.013		0.183***		0.200***
		(0.077)		(0.061)		(0.056)
FS: Δ log emp 16-64 (adj)						
Current Bartik shock		0.295***		0.253*		0.244*
		(0.112)		(0.135)		(0.138)
Lagged Bartik shock		0.057		0.414***		0.347***
		(0.084)		(0.109)		(0.101)
Current Bartik shock (adj)		1.255***		0.750***		0.622***
		(0.128)		(0.161)		(0.182)
Lagged Bartik shock (adj)		-0.659***		-0.862***		-1.019***
		(0.119)		(0.139)		(0.133)
FS: Lagged emp rate 16-64 (adj)						
Current Bartik shock		0.035		0.013		0.075**
		(0.059)		(0.038)		(0.032)
Lagged Bartik shock		-0.018		-0.137***		-0.079**
		(0.062)		(0.045)		(0.034)
Current Bartik shock (adj)		0.05		0.258***		0.199***
		(0.106)		(0.060)		(0.053)
Lagged Bartik shock (adj)		0.529***		0.463***		0.363***
		(0.077)		(0.050)		(0.037)
Observations	2,876	2,876	2,876	2,876	2,157	2,157

This table reports OLS and IV estimates of the elasticity of the share of employed residents working locally to employment shocks, as well as the first stage estimates, across 722 CZs and four (decadal) time periods beginning in 1970 (when our commuting data begins). In columns 1, 3 and 5, the right-hand side of the empirical specification is identical to that of Table 2, based on the population response equation (15). In the remaining columns, we also control for the change in the log total employment in all adjacent CZs and the lagged log employment rate across those CZs. To this end, we include two further instruments: the current and lagged (employment-weighted) average Bartik shift-share across adjacent CZs. We exclude from our sample those 3 CZs which are islands (with no adjacent CZs): these are Martha's Vineyard (MA), Nantucket (MA) and San Juan Islands (WA). See the notes under Table 2 for further details on the empirical specification. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1. *** p<0.01, ** p<0.05, * p<0.1.

Table 12: Robustness checks for population response

PANEL A: WEIGHTED BY LAGGED LOCAL POP SHARE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ log emp 16-64	0.711***	0.703***	0.661***	0.670***	0.688***	0.684***	0.664***	0.819***	0.813***	0.715***
	(0.030)	(0.033)	(0.029)	(0.028)	(0.022)	(0.022)	(0.026)	(0.020)	(0.037)	(0.030)
Lagged log emp rate 16-64	0.268***	0.191***	0.362***	0.401***	0.473***	0.457***	0.445***		1.037***	0.796***
	(0.053)	(0.044)	(0.048)	(0.051)	(0.054)	(0.048)	(0.051)		(0.156)	(0.101)
Migrant shift-share		0.228***	0.01	0.057	0.101	0.073	0.056	0.093***	0.412***	0.316***
		(0.058)	(0.058)	(0.062)	(0.063)	(0.063)	(0.068)	(0.031)	(0.124)	(0.089)
Max temp January			0.236***	0.246***	0.207***	0.212***	0.141***	0.016		
			(0.022)	(0.024)	(0.025)	(0.025)	(0.036)	(0.031)		
Max temp July			-0.095**	-0.129***	-0.171***	-0.165***	-0.093	0.002		
			(0.042)	(0.048)	(0.047)	(0.046)	(0.072)	(0.063)		
Mean humidity July			-0.080***	-0.072***	-0.033**	-0.034**	0.053**	0.081***		
			(0.016)	(0.017)	(0.015)	(0.014)	(0.025)	(0.023)		
Coastline dummy				-0.008**	-0.007*	-0.009**	0.01	0.007		
				(0.004)	(0.004)	(0.003)	(0.007)	(0.006)		
Log pop density 1900					-0.007***	-0.005***	-0.012***	-0.010***		
					(0.001)	(0.001)	(0.003)	(0.003)		
Log distance to closest CZ						0.015**	0.015	-0.01		
						(0.007)	(0.013)	(0.012)		
Year effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Amenities x year effects	No	No	No	No	No	No	Yes	Yes	Yes	Yes
CZ fixed effects	No	No	No	No	No	No	No	No	Yes	No
First-differenced spec	No	No	No	No	No	No	No	No	No	Yes
Observations	4,332	4,332	4,332	4,332	4,332	4,332	4,332	4,332	4,332	3,610

PANEL B: UNWEIGHTED

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ log emp 16-64	0.646***	0.627***	0.601***	0.616***	0.615***	0.604***	0.621***	0.739***	0.656***	0.674***
	(0.024)	(0.024)	(0.021)	(0.020)	(0.021)	(0.021)	(0.026)	(0.023)	(0.028)	(0.024)
Lagged log emp rate 16-64	0.580***	0.510***	0.417***	0.449***	0.444***	0.416***	0.383***		0.572***	0.691***
	(0.052)	(0.046)	(0.035)	(0.038)	(0.040)	(0.039)	(0.036)		(0.083)	(0.076)
Migrant shift-share		0.577***	0.199***	0.237***	0.236***	0.182**	0.084	0.066*	-0.008	0.062
		(0.074)	(0.077)	(0.079)	(0.078)	(0.075)	(0.074)	(0.037)	(0.088)	(0.088)
Max temp January			0.365***	0.380***	0.379***	0.370***	0.245***	0.139***		
			(0.020)	(0.021)	(0.020)	(0.020)	(0.026)	(0.023)		
Max temp July			-0.494***	-0.529***	-0.528***	-0.488***	-0.338***	-0.221***		
			(0.039)	(0.040)	(0.039)	(0.038)	(0.058)	(0.051)		
Mean humidity July			-0.028**	-0.017	-0.021	-0.007	0.058*	0.075***		
			(0.013)	(0.013)	(0.018)	(0.018)	(0.030)	(0.026)		
Coastline dummy				-0.015***	-0.015***	-0.018***	0.014*	0.01		
				(0.005)	(0.005)	(0.005)	(0.008)	(0.007)		
Log pop density 1900					0.001	0.003*	0.000	-0.001		
					(0.002)	(0.002)	(0.003)	(0.003)		
Log distance to closest CZ						0.037***	0.046***	0.036***		
						(0.007)	(0.012)	(0.011)		
Year effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Amenities x year effects	No	No	No	No	No	No	Yes	Yes	Yes	Yes
CZ fixed effects	No	No	No	No	No	No	No	No	Yes	No
First-differenced spec	No	No	No	No	No	No	No	No	No	Yes
Observations	4,332	4,332	4,332	4,332	4,332	4,332	4,332	4,332	4,332	3,610

This table tests robustness of our IV estimates of the population response in Table 2, to contemporaneous employment growth and the lagged employment rate. As before, our sample covers the 722 CZs and six (decadal) time periods. First, we test robustness of our estimates to the weighting of observations: Panel A weights observations of the lagged population share, and Panel B applies no weighting. And second, we test robustness to the inclusion of progessively more controls. Columns 1-7 do not condition on CZ effects, and the final two columns report the fixed effect and first differenced specifications. Notice that columns 7, 9 and 10 in Panel A are identical to columns 4, 5 and 6 respectively of Table 2 above. In column 8, we show what happens if the lagged ECM term (and its lagged Bartik instrument) are omitted. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

Table 13: First stage estimates by demographic group

	Δ log emp 16-64			Lagged emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
College-graduates						
Current Bartik shock	0.860***	0.715***	0.566***	-0.170***	-0.115**	-0.023
	(0.091)	(0.093)	(0.124)	(0.048)	(0.057)	(0.048)
Lagged Bartik shock	-0.294***	-0.537***	-0.511***	0.166***	0.171***	0.162***
	(0.083)	(0.100)	(0.113)	(0.043)	(0.056)	(0.040)
Non-graduates						
Current Bartik shock	1.091***	1.137***	1.051***	-0.034	-0.164***	-0.059*
	(0.067)	(0.076)	(0.081)	(0.042)	(0.036)	(0.034)
Lagged Bartik shock	-0.129**	-0.145**	-0.219**	0.535***	0.222***	0.267***
	(0.059)	(0.065)	(0.087)	(0.053)	(0.037)	(0.026)
Men						
Current Bartik shock	1.246***	1.221***	1.197***	-0.196***	-0.294***	-0.183***
	(0.073)	(0.082)	(0.087)	(0.033)	(0.037)	(0.027)
Lagged Bartik shock	-0.026	-0.135**	-0.142	0.451***	0.188***	0.192***
	(0.063)	(0.066)	(0.094)	(0.040)	(0.037)	(0.027)
Women						
Current Bartik shock	0.792***	0.725***	0.626***	0.644***	0.416***	0.321***
	(0.064)	(0.064)	(0.076)	(0.086)	(0.069)	(0.068)
Lagged Bartik shock	-0.076	-0.188***	-0.314***	0.634***	0.230***	0.270***
	(0.059)	(0.069)	(0.084)	(0.083)	(0.045)	(0.031)
16-24s						
Current Bartik shock	1.601***	1.616***	1.639***	-0.102	-0.312***	-0.195***
	(0.098)	(0.114)	(0.133)	(0.070)	(0.055)	(0.052)
Lagged Bartik shock	-0.191**	-0.234**	-0.286**	0.848***	0.305***	0.379***
	(0.084)	(0.113)	(0.137)	(0.087)	(0.074)	(0.049)
25-44s						
Current Bartik shock	1.238***	1.197***	1.023***	0.05	-0.05	0.027
	(0.075)	(0.085)	(0.097)	(0.034)	(0.030)	(0.038)
Lagged Bartik shock	-0.124*	-0.245***	-0.292***	0.377***	0.182***	0.219***
	(0.069)	(0.070)	(0.101)	(0.040)	(0.031)	(0.025)
45-64s						
Current Bartik shock	0.789***	0.740***	0.727***	-0.012	-0.199***	-0.088***
	(0.069)	(0.067)	(0.080)	(0.041)	(0.049)	(0.030)
Lagged Bartik shock	0.143***	0.023	-0.02	0.507***	0.230***	0.245***
	(0.054)	(0.058)	(0.063)	(0.056)	(0.031)	(0.026)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

This table reports first stage estimates from the group-specific IV regressions in Tables 5 and 6, based on the population response equation (15). The endogenous variables in each case are the local change in log employment and the lagged log employment rate, within the specified demographic group. See the notes accompanying Tables 5 and 6 for further details on the empirical specification. The sample covers 722 CZs and six (decadal) time periods. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share. *** p<0.01, ** p<0.05, * p<0.1.

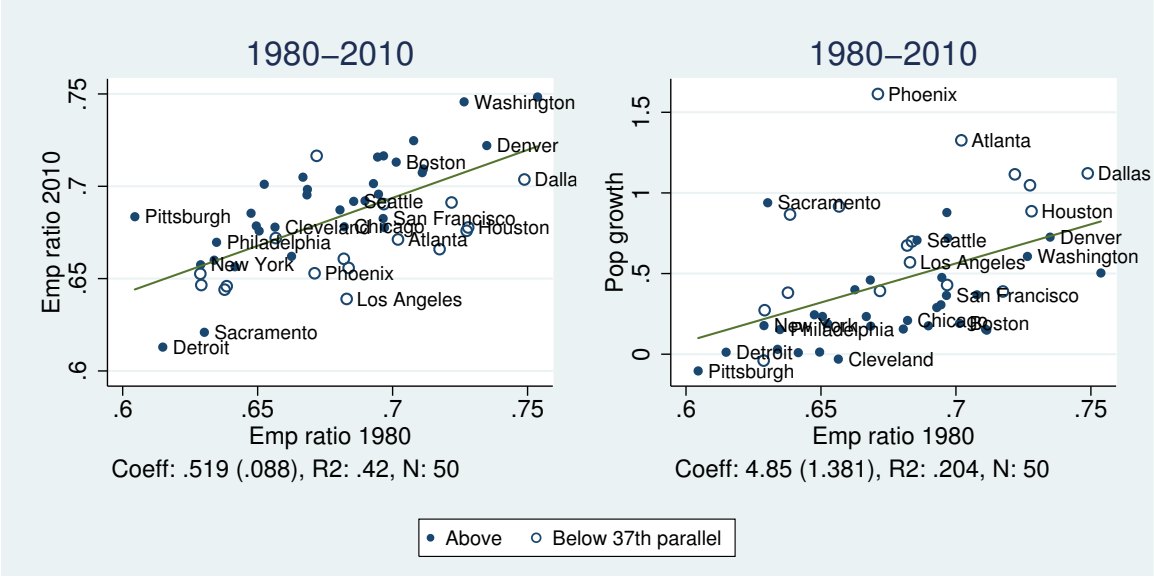


Figure 1: Persistence in employment ratio and population response

Note: Data-points denote commuting zones (CZs). Sample is restricted to the 50 largest commuting zones in 1980, and divided into CZs above and below the 37th parallel (where the latter corresponds to the Sunbelt region). Employment rate and population growth are estimated for individuals aged 16-64.

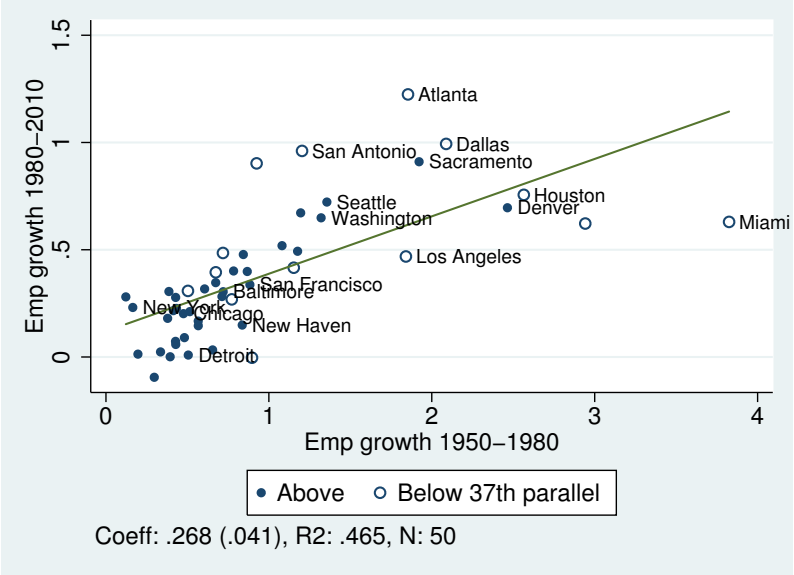


Figure 2: Persistence in local employment growth

Note: Data-points denote commuting zones. Sample is restricted to the 50 largest commuting zones in 1950. Employment growth is estimated for individuals aged 16-64.

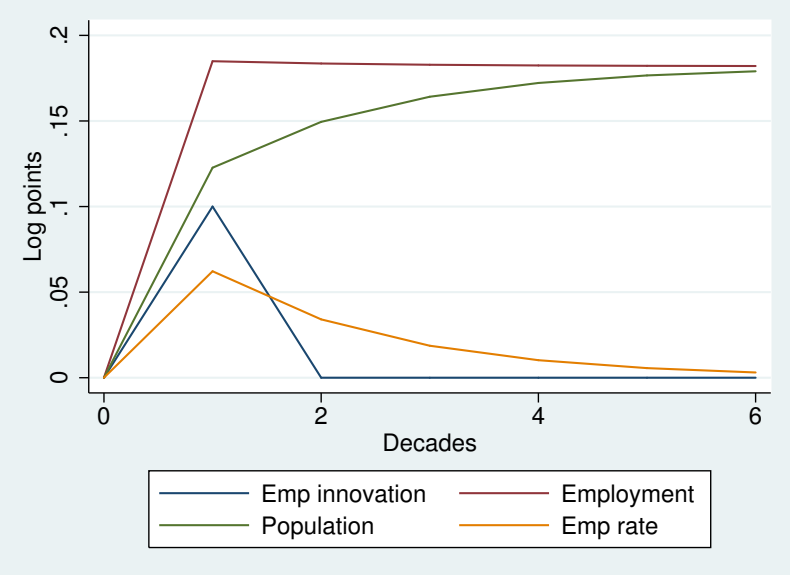


Figure 3: Impulse response function

Note: This figure illustrates the impulse response following an innovation of 0.1 log points in z_{rt} for some area r at $t = 1$, from an initial position of a steady-state equilibrium. The response is based on our preferred estimates of the β and γ parameters.

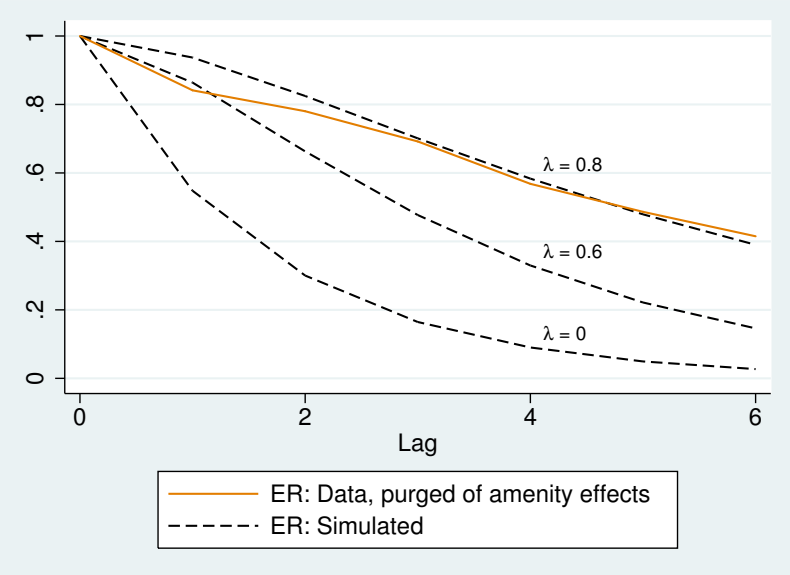


Figure 4: ACFs of employment rate: data and simulated for different λ

Note: This figure illustrates the observed and simulated persistence of the log employment rate. The orange line is the time-demeaned ACF of the log employment rate in the data, purged of observable amenity effects (the 10th row in Table 1). The dashed lines are the simulated ACFs for different values of λ in equation (24), given our preferred estimates of the β and γ parameters.

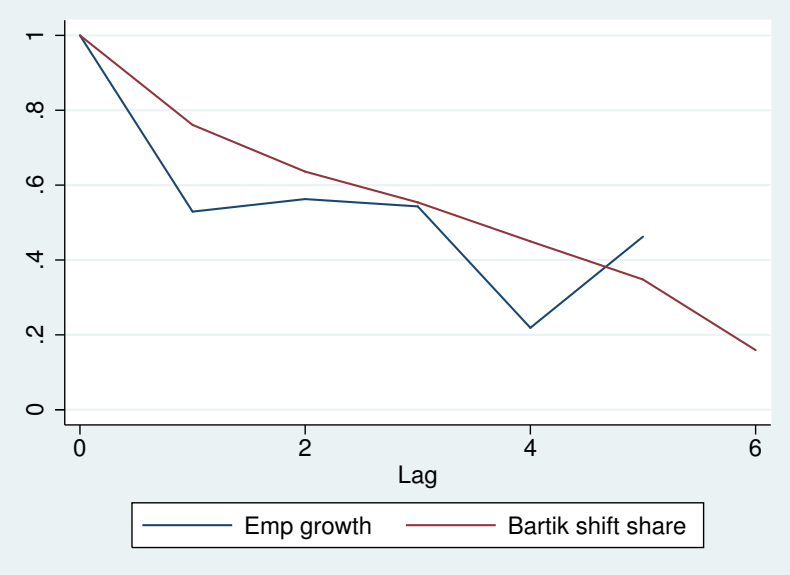


Figure 5: ACFs for local employment growth and Bartik series

Note: This figure illustrates the time-demeaned ACFs of the change in log CZ employment and the Bartik shift-share described in equation (16). The estimation of the ACFs is described in the notes accompanying Table 1. Our panel for the Bartik shift-share variable is one (decadal) period longer than for local employment growth, with the former including the 1940s. As described in the Appendix B, the employment counts of the 1940 census are not comparable with those of later years.

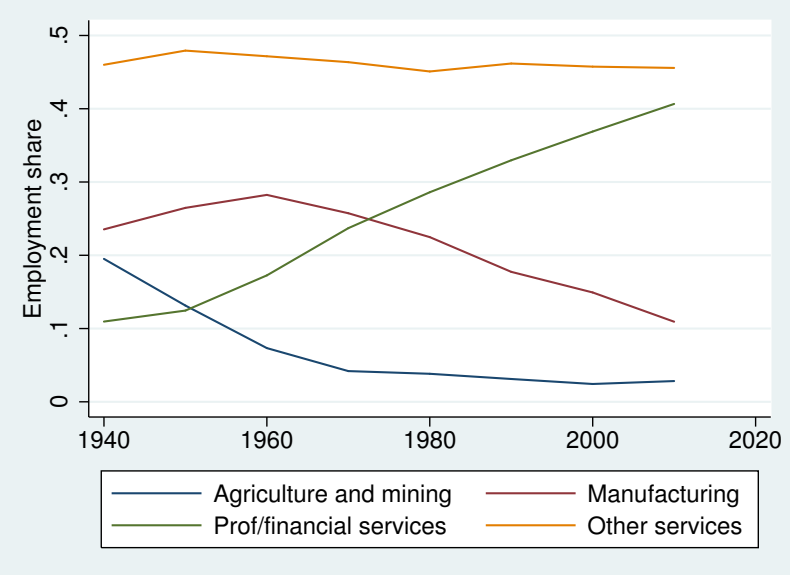


Figure 6: Employment shares by industry

Note: Professional and financial services include all medical, legal, engineering, architectural, accounting, advertising, financial, insurance and real estate services.

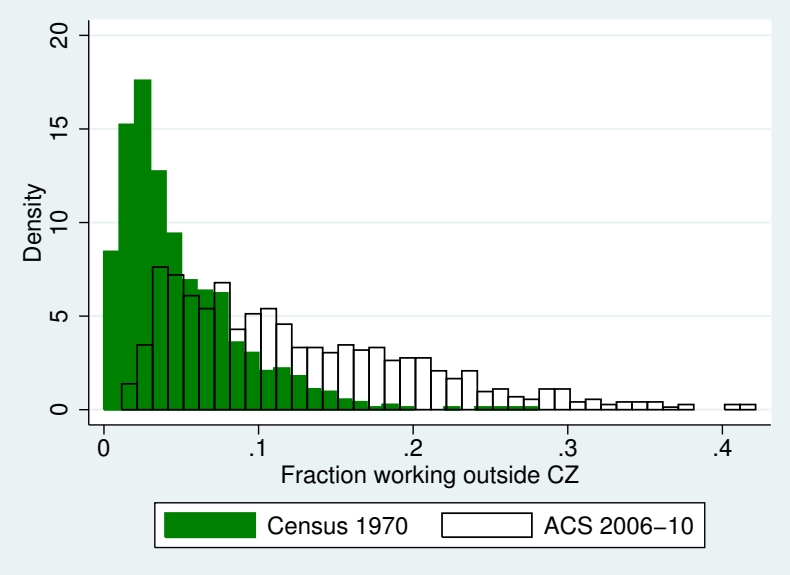


Figure 7: Histograms (across CZs) of fraction employed outside CZ: 1970 and 2006-10

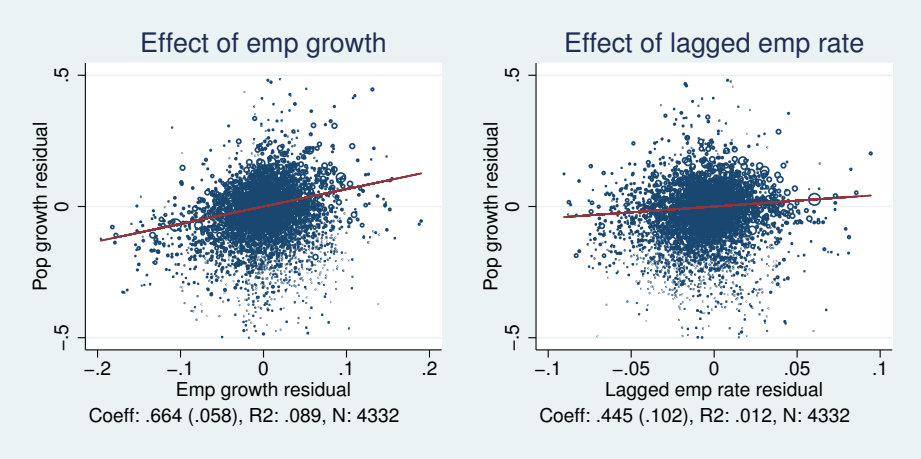


Figure 8: Graphical illustration of 2SLS coefficients

Note: Following the Frisch-Waugh theorem, this figure graphically depicts the 2SLS coefficients of the basic specification (column 4, Table 2; column 7, Table 12). The first panel illustrates the coefficient on employment growth, and the second panel the coefficient on the lagged employment rate. We construct these graphs in two steps. We use the first stage regressions to predict a full set of values for each endogenous variable. Then, on the y-axis of the first panel, we plot the residuals from a regression of population growth on the predicted lagged employment growth, together with all exogenous variables. On the x-axis, we plot the residuals of a regression of employment growth on the same set of explanatory variables. In the second panel, we repeat the exercise for the lagged employment rate. Data-points denote CZ-year observations, with size corresponding to lagged population. Standard errors for the best-fit slope are heteroskedasticity-consistent and clustered by CZ. Notice the standard errors of the best-fit slopes do not correspond to those in Table \ref{tab: Basic estimates}); this is because this naive estimator does not account for the standard error in the first stage. A small number of outlying data points (no more than 30 minor CZs in each panel) have been excluded because of our choice of axis range.

Bibliography

Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2014): "Import Competition and the Great US Employment Sag of the 2000s," Tech. rep.

Albouy, D. (2012): "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas," Tech. rep.

Amior, M. (2014): "Why are Skilled Workers More Mobile Geographically? The Role of the Job Surplus," Tech. rep., unpublished.

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *The American Economic Review*, 103, 1553–1597.

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

Autor, D. H. and M. G. Duggan (2003): "The rise in the disability rolls and the decline in unemployment," *The Quarterly Journal of Economics*, 157–205.

Bartik, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?*, WE Upjohn Institute for Employment Research.

Bayer, C. and F. Juessen (2012): "On the Dynamics of Interstate Migration: Migration Costs and Self-Selection," *Review of Economic Dynamics*, 15, 377–401.

Beaudry, P., D. A. Green, and B. M. Sand (2012): "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory," *Econometrica*, 80, 1063–1104.

——— (2014a): "In Search of Labor Demand," National Bureau of Economic Research Working Paper No. 20568.

——— (2014b): "Spatial Equilibrium with Unemployment and Wage Bargaining: Theory and Estimation," *Journal of Urban Economics*, 79, 2–19.

Bertola, G. and A. Ichino (1995): "Wage inequality and unemployment: United States versus Europe," in *NBER Macroeconomics Annual 1995, Volume 10*, MIT Press, 13–66.

Beyer, R. C. and F. Smets (2014): "Labour Market Adjustments and Migration in Europe and the United States: How Different?" .

Blanchard, O. and L. F. Katz (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 23, 1–76.

Blanchflower, D. G. and A. J. Oswald (1994): *The wage curve*, MIT Press.

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

Card, D. (2001): "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, 19, 22–64.

Center, M. P. (2011): "National Historical Geographic Information System: Version 2.0," Tech. rep., University of Minnesota.

Dao, M., D. Furceri, and M. P. Loungani (2014): *Regional Labor Market Adjustments in the United States*, 14-211, International Monetary Fund.

Davis, M. A. and F. Ortalo-Magne (2011): "Household expenditures, wages, rents," *Review of Economic Dynamics*, 14, 248–261.

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

Diamond, R. (2013): "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," http://web.stanford.edu/∼diamondr.

Eichengreen, B. (1993): *Policy Issues in the Operation of Currency Unions*, Cambridge University Press, Cambridge, chap. Labor markets and European monetary unification, 130–162.

Gallin, J. (2008): "The Long-Run Relationship Between House Prices and Rents," *Real Estate Economics*, 36, 635–658.

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

Glaeser, E. L. (2005): "Urban colossus: why is New York America's largest city?" Tech. rep., National Bureau of Economic Research.

Glaeser, E. L. and J. D. Gottlieb (2009): "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States," *Journal of Economic Literature*, 47, 983–1028.

Glaeser, E. L. and J. Gyourko (2005): "Urban Decline and Durable Housing," *Journal of Political Economy*, 113, 345–000.

Glaeser, E. L., J. Gyourko, and R. E. Saks (2006): "Urban growth and housing supply," *J Econ Geogr*, 6, 71–89.

Glaeser, E. L., J. Kolko, and A. Saiz (2001): "Consumer city," *Journal of economic geography*, 1, 27–50.

Glaeser, E. L. and G. A. Ponzetto (2007): "Did the Death of Distance Hurt Detroit and Help New York?" Tech. rep., National Bureau of Economic Research.

Glaeser, E. L. and C. Redlick (2009): "Social capital and urban growth," *International Regional Science Review*.

Glaeser, E. L., M. Resseger, and K. Tobio (2009): "INEQUALITY IN CITIES*," *Journal of Regional Science*, 49, 617–646.

Glaeser, E. L., J. Scheinkman, and A. Shleifer (1995): "Economic growth in a cross-section of cities," *Journal of Monetary Economics*, 36, 117–143.

Glaeser, E. L. and J. Shapiro (2001): "Is there a new urbanism? The growth of US cities in the 1990s," Tech. rep., National Bureau of Economic Research.

Glaeser, E. L. and K. Tobio (2007): "The rise of the sunbelt," Tech. rep., National Bureau of Economic Research.

Gonzalez, A. (1998): "Mexican enclaves and the price of culture," *Journal of Urban Economics*, 43, 273–291.

Hughes, G. and B. McCormick (1987): "Housing markets, unemployment and labour market flexibility in the UK," *European Economic Review*, 31, 615–641.

Kaplan, G. and S. Schulhofer-Wohl (2012): "Understanding the Long-Run Decline in Interstate Migration," Tech. rep., National Bureau of Economic Research Working Paper No. 18507.

Kennan, J. (2013): "Higher Education Subsidies and Human Capital Mobility," Tech. rep.

Kennan, J. and J. R. Walker (2011): "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79, 211–251.

Kline, P. and E. Moretti (2013): "Place Based Policies with Unemployment," *The American Economic Review*, 103, 238–243.

Koo, J., K. R. Phillips, and F. D. Sigalla (2000): "Measuring regional cost of living," *Journal of Business & Economic Statistics*, 18, 127–136.

McGranahan, D. A. (1999): "Natural Amenities Drive Rural Population Change," Tech. rep., Agricultural Economic Report No. 781, U.S. Department of Agriculture.

Molloy, R., C. L. Smith, and A. Wozniak (2011): "Internal Migration in the United States," *Journal of Economic Perspectives*, 25, 173–96.

Moretti, E. (2004): "Human capital externalities in cities," *Handbook of regional and urban economics*, 4, 2243–2291.

——— (2010): "Local multipliers," *The American Economic Review*, 373–377.

——— (2011): "Local Labor Markets," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, New York: Elsevier, vol. 4B, 1237–1313.

Munshi, K. (2003): "Networks in the modern economy: Mexican migrants in the US labor market," *The Quarterly Journal of Economics*, 118, 549–599.

Nickell, S. (1981): "Biases in dynamic models with fixed effects," *Econometrica: Journal of the Econometric Society*, 1417–1426.

Notowidigdo, M. J. (2011): "The Incidence of Local Labor Demand Shocks," Tech. rep., National Bureau of Economic Research Working Paper No. 17167.

Obstfeld, M. and G. Peri (1998): "Regional non-adjustment and fiscal policy," *Economic Policy*, 207–259.

Oi, W. Y. (1996): "The welfare implications of invention," in *The Economics of New Goods*, University of Chicago Press, 109–142.

Overman, H. G. and D. Puga (2002): "Unemployment clusters across Europe's regions and countries," *Economic policy*, 17, 115–148.

Phillips, K. R. and C. Daly (2010): "Improving the ACCRA US regional cost of living index," *Journal of Economic and Social Measurement*, 35, 33–42.

Rappaport, J. (2007): "Moving to nice weather," *Regional Science and Urban Economics*, 37, 375–398.

——— (2012): "Why does unemployment differ persistently across metro areas?" *Federal Reserve Bank of Kansas City, Economic Review*, 97, 5–35.

Rappaport, J. and J. D. Sachs (2003): "The United States as a coastal nation," *Journal of Economic growth*, 8, 5–46.

Roback, J. (1982): "Wages, Rents, and the Quality of Life," *The Journal of Political Economy*, 1257–1278.

Rosen, S. (1979): *Current Issues in Urban Economics*, Baltimore: Johns Hopkins University Press, chap. Wage-based indexes of urban quality of life, 74–104.

Saiz, A. (2010): "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

Tolbert, C. M. and M. Sizer (1996): "U.S. Commuting Zones and Labor Market Areas: A 1990 Update," Tech. rep., Economic Research Service Staff Paper No. 9614.

Topel, R. H. (1986): "Local Labor Markets," *The Journal of Political Economy*, 94, S111–S143.

Wozniak, A. (2010): "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources*, 45, 944–970.

Yagan, D. (2014): "Moving to opportunity? Migratory insurance over the Great Recession," http://eml.berkeley.edu/∼yagan/.