

The Persistence of Local Joblessness

Michael Amior and Alan Manning*

January 2018

Abstract

Differences in employment-population ratios across US commuting zones have persisted for many decades. We claim these disparities represent real gaps in economic opportunity for individuals of fixed characteristics. These gaps persist despite a strong migratory response, and we attribute this to high persistence in labor demand shocks. These trends generate a “race” between local employment and population: population always lags behind employment, yielding persistent deviations in employment rates. Methodologically, we argue the employment rate can serve as a sufficient statistic for local well-being; and we model population and employment dynamics using an error correction mechanism, which explicitly allows for disequilibrium.

1 Introduction

It is well known that local joblessness is very persistent: see Kline and Moretti (2013) on the US, Overman and Puga (2002) on Europe, and OECD (2005*a*) for cross-country comparisons. This is illustrated in the first panel of Figure 1, which compares employment-population ratios (from here on, “employment rates”) in 1980 and 2010 among 16-64s, for the 50 largest US commuting zones. The correlation is 0.42. We show later this persistence for both men and women and in both labor force participation and unemployment rates (see Online Appendix E for further graphical illustrations).

The persistence might naturally be interpreted as an equilibrium phenomenon, driven by local variation in demographic composition or compensated by local amenities. In this view,

*Amior: Hebrew University of Jerusalem, Mount Scopus, Jerusalem 91905, Israel, and Centre for Economic Performance; michael.amior@mail.huji.ac.il. Manning: London School of Economics, Houghton Street, London WC2A 2AE, UK, and Centre for Economic Performance; a.manning@lse.ac.uk. We are grateful to the editor, referees and seminar participants at the Centre for Economic Performance, the Spatial Economics Research Centre, the SOLE-EALE conference, the NBER Summer Institute and the University of Cambridge, and to Jack DeWaard, Joe Grover, Kin Koerber, Enrico Moretti and Jordan Rappaport for sharing data. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

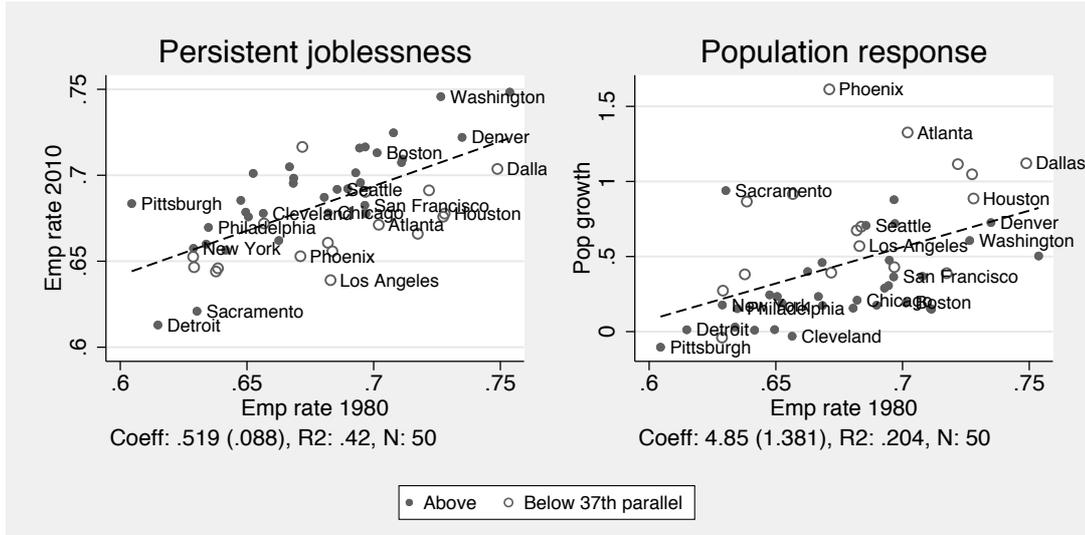


Figure 1: Persistence in employment-population ratios and population response

Data-points denote commuting zones (CZs). Sample is restricted to the 50 largest CZs in 1980, for individuals aged 16-64, and divided into CZs above and below the 37th parallel (i.e. the Sun Belt).

utility is equalized across areas (at least for marginal residents) and we would therefore not expect any systematic relationship between employment rates and population growth. However, the second panel of Figure 1 shows a strong population response: those areas with the highest employment rates in 1980 grew by over 50 percentage points more in the subsequent three decades than those with the lowest.¹ This suggests that persistent joblessness is not an equilibrium phenomenon and that migration contributes to local adjustment - as Blanchard and Katz (1992) show. But given this population response, it is surprising that jobless rates are so persistent. The central aim of this paper is to explain how persistent joblessness and a strong migratory response can coexist.

We claim that large persistence in local labor demand shocks can resolve this puzzle. As Figure 2 shows, employment growth between 1950 and 1980 is strongly correlated with growth between 1980 and 2010.² We argue this generates a “race” between local employment and population, a geographical analogue of the more famous race from Tinbergen (1974) and Goldin and Katz (2008) between technology (on the demand side) and skills (supply). While local population does respond strongly to demand shocks, it always lags behind employment. This results in fairly static local deviations in jobless rates, despite

¹These population responses to initial employment rates were previously documented by Glaeser, Scheinkman and Shleifer (1995) and Glaeser and Shapiro (2001) using similar data.

²Similar patterns have previously been documented by Blanchard and Katz (1992) and Dao, Furceri and Loungani (2017).

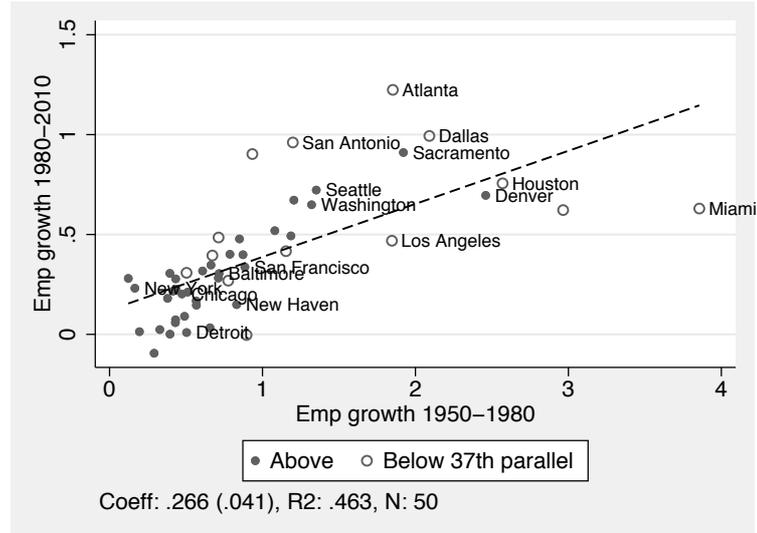


Figure 2: Persistence in local employment growth

Data-points denote CZs. Sample is restricted to 50 largest commuting zones in 1950, for individuals aged 16-64.

large changes in employment and population.

In Section 2, we use a variant of the classic Rosen-Roback framework to investigate the dynamics of population and employment (Rosen, 1979; Roback, 1982). We begin by characterizing local equilibrium for a fixed population. We modify the standard model by including a classical labor supply curve or “wage curve” (Blanchflower and Oswald, 1994), so labor supply is not completely inelastic. This means that local demand shocks affect employment rates and not just wages, which is consistent with existing evidence (Blanchard and Katz, 1992; Beaudry, Green and Sand, 2014b). One contribution of the paper is to show how the employment rate can then serve as a “sufficient statistic” for local economic well-being, 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 their empirical model. This change in focus from wages to employment rates has practical advantages: employment rates are easier to measure than real 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 adjustment.

We combine our local equilibrium model for fixed population with a dynamic migration equation: workers move to areas offering higher utility, but this process takes time. This yields an error correction model (ECM), where changes in log local population are influenced by (i) changes in log local employment (driven by demand shocks) and (ii) the lagged log

local employment rate (representing the extent of disequilibrium). We estimate this model on decadal census data going back to 1950 for 722 commuting zones (CZs).

This ECM approach naturally encapsulates the “race” between population and employment. Many existing studies of decadal population growth control for contemporaneous shocks, but they typically omit the disequilibrium term (e.g. Bound and Holzer, 2000; Glaeser, Gyourko and Saks, 2006; Notowidigdo, 2011; Autor, Dorn and Hanson, 2013; Beaudry, Green and Sand, 2014*b*; Diamond, 2016). This can be problematic if the two are correlated, as they will be if demand shocks are persistent. There is already evidence that these long run dynamics matter. For example, Hornbeck (2012) shows that population adjustment away from the Great Plains, following the 1930s Dust Bowl crisis, continued through the 1950s. And the fact that local shocks have long-lasting effects on labor force participation (see Bound and Holzer, 2000; Black, Daniel and Sanders, 2002; Autor and Duggan, 2003; Autor, Dorn and Hanson, 2013; Yagan, 2017) is suggestive of sluggish adjustment. Partridge et al. (2012) and Beaudry, Green and Sand (2014*b*) make the explicit claim that the system may not be close to spatial equilibrium. And Jaeger, Ruist and Stuhler (2017) have emphasized the long run dynamics in their analysis of local adjustment to foreign migration.

Section 3 describes the data and estimates local persistence in the employment rate. This is remarkably strong over many decades, for both men and women, and both before and after 1980 - though it is weaker among better educated workers and labor force participants. Critically, we argue this persistence reflects real differences in employment opportunities for individuals of fixed characteristics. In particular, it cannot be explained by individuals trading off employment opportunities against permanent amenities such as climate (see e.g. Roback, 1982) or differences in demographic composition. Adjusting local employment rates for a detailed set of amenities and individual characteristics does little to the estimated persistence.

Section 4 presents estimates of our model for population growth. We instrument contemporaneous employment growth with an industry shift-share (following Bartik, 1991) and the lagged employment rate with the lagged shift-share. The model fits the data well: our preferred estimates imply that, over one decade, population corrects for 40 percent of the initial deviation in the local employment rate - a large response, but more sluggish than in Blanchard and Katz (1992). We also find some heterogeneity in these responses: the college graduate population adjusts more quickly than the non-graduate, the prime-aged population more quickly than older workers, and the labor force more quickly than the general population. There is some (mixed) evidence that the population response fell over time, which may

reflect declining gross migration rates (see e.g. Molloy, Smith and Wozniak, 2011; Partridge et al., 2012; Beyer and Smets, 2015; Dao, Furceri and Loungani, 2017).

The estimates discussed so far take employment as given. But to derive the dynamic response to a local demand shock, we must recognize that employment is itself endogenous. In Section 5, we estimate an ECM for employment growth. Employment responds negatively to the initial employment rate (which hastens adjustment) and positively to population. Based on our preferred estimates of the two ECMs, we predict the employment rate follows an AR(1) process with a decadal persistence of 0.68. Though this number is substantial, it is clearly insufficient to explain the degree of persistence (over 30 years) illustrated 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 0.5 and 0.8 (in the demand innovation process) is needed to achieve this. This is a high number, but the persistence in our demand instrument (the industry shift-share) falls in the same range; and this reassures us that these numbers are plausible.

As well as providing an explanation for the persistence in joblessness (which we have argued represents real differences in utility for individuals of fixed characteristics), the paper makes two methodological contributions: first, that the employment rate can serve as a sufficient statistic for local economic well-being; and second, that an error correction mechanism that explicitly allows for disequilibrium is a useful way of modeling population and employment dynamics. The ECM captures the idea of a race between local employment and population in the face of persistent demand shocks. Although migration helps reduce these disparities, serial correlation in local demand growth ensures it is insufficient to equalize economic opportunity across areas.

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), and the survey by Spring, Tolnay and Crowder (2016) for a broader perspective on these issues which spans many disciplines. First, we describe the short run equilibrium with local population fixed. And then, we show how this can be combined with a dynamic migration equation to yield our ECM model.

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, P_r , in each area given by:

$$P_r = Q(P, P_r^h) \quad (1)$$

where $Q()$ is a price index. We assume here that individuals are homogeneous ex ante (see Online Appendix A.6 for a discussion of alternative assumptions), though we do seek to convince that our results are not driven by omitted skill differences. Total population is given by L_r and employment by N_r . In the standard Rosen-Roback framework, employment and population are synonymous; but this is an important distinction to make in our model. Individuals earn W_r in employment and B when out of work, e.g. through unemployment insurance or disability benefits. Labor income is taxed at rate τ (see e.g. Albouy, 2008). Total income in an area is then given by³:

$$Y_r = (1 - \tau) W_r N_r + B(L_r - N_r) \quad (2)$$

Using (1) and basic consumer theory, housing demand can then be written in log-linearized form as:

$$\begin{aligned} h_r^d &= y_r - p + \epsilon^{hd} (p_r^h - p) \\ &= w_r - p + l_r + \kappa (n_r - l_r) + \epsilon^{hd} (p_r^h - p) \end{aligned} \quad (3)$$

where lower case variables denote logs, h_r^d is the log of housing demand, ϵ^{hd} the price elasticity of housing demand, and κ its elasticity with respect to the employment rate. The final line substitutes a log-linear approximation of (2).

We now turn to production. We assume here that housing production does not depend on local labor, though Online Appendix A.1 shows this does not affect the main results. Housing supply can then be written as:

$$h_r^s = \epsilon_r^{hs} (p_r^h - p) \quad (4)$$

where h_r^s is the log of housing supply and ϵ_r^{hs} the price elasticity of housing supply, which we allow to vary by area in line with Saiz (2010). We also assume production of the traded

³This implicitly assumes that everyone rents their housing, which rules out the possibility of wealth effects if individuals own housing equity. We discuss the theoretical and empirical implications of these wealth effects in Section 4.3.

good does not depend on land. Consequently, the demand for labor in area r can be written as:

$$n_r^d = \epsilon^{nd} (w_r - p) + z_r^d \quad (5)$$

where n_r^d is the log of labor demand, ϵ^{nd} the wage elasticity of labor demand, and z_r^d an area-specific shifter that will be the source of local shocks. Implicitly, we assume capital is supplied perfectly elastically at a regionally invariant cost; one could alter this assumption and model capital more explicitly. To close the model, we need a labor supply curve or its non-competitive equivalent. In contrast to the standard Rosen-Roback framework, we will assume there is some elasticity in this relationship:

$$n_r^s = l_r + \epsilon^{ns} (w_r - p_r) + z_r^s \quad (6)$$

where n_r^s is the log of labor supply, ϵ^{ns} the wage elasticity of labor supply, and z_r^s an area-specific deviation in the labor supply curve. Assuming $\epsilon^{ns} > 0$, there is a positive relationship between an area's real consumer wage and its employment rate. (6) can be interpreted as an elastic labor supply curve if the labor market is competitive, or as a "wage curve" (Blanchflower and Oswald, 1994) in some other labor market model.

Consider a short run equilibrium with local population fixed. If labor is inelastically supplied, an increase in demand z_r^d is fully manifested in wages. From (3), this raises housing demand and prices. Prices grow more if housing supply is less elastic, but this will not generally offset the growth in wages; so the real consumer wage increases overall. If labor supply is somewhat elastic, the effect on real wages is more muted as the employment rate also grows.

We define expected utility in area r as:

$$u_r = \sigma (n_r - l_r) + (w_r - p_r) + a_r \quad (7)$$

where a_r is the value of local amenities. Utility is a function of the employment rate and the real consumer wage. But, it is difficult to empirically disentangle the effect of employment and wages.⁴ In particular, a labor demand shock in this model will not yield independent

⁴See Beaudry, Green and Sand (2014b) for one attempt to do so. Interestingly, they show the local population response to employment is larger and more robust than to wages or prices. They suggest this may be because employment transitions are discrete shocks with larger implications for individual income. In Online Appendix J, we attempt a similar strategy to Beaudry, Green and Sand to simultaneously identify the effect of both employment and wages; and we also find the employment effect is much more robust.

variation in these two variables. In these circumstances, either the employment rate or real consumer wage alone can serve as a one-dimensional measure of local labor market conditions: the other variable can be eliminated in (7) using the labor supply curve (6). In particular, substituting for the real consumer wage gives:

$$u_r = \left(\sigma + \frac{1}{\epsilon^{ns}} \right) (n_r - l_r) + a_r - \frac{1}{\epsilon^{ns}} z_s^r \quad (8)$$

so utility is reduced to the employment rate and an area fixed effect, which is an amalgam of the amenity effect a_r and labor supply shifter z_s^r . Blanchard and Katz (1992) implicitly perform a similar transformation in deriving a VAR model in local quantities (employment growth, employment rate and participation rate). One does lose information in reducing utility to one dimension; 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 areas⁵, though this will influence the extent to which labor demand shocks affect the employment rate.⁶

One could perform a similar substitution to eliminate the employment rate from (7) and write everything in terms of the real consumer wage as most of the literature does. There is nothing to choose between the approaches theoretically, except in the extreme cases where labor supply is perfectly elastic or inelastic. Practically though, local wage deflators are available for shorter periods and fewer geographies than employment rates and are notoriously difficult to estimate in any case.⁷ Also, employment is measured in the same units as population, and this allows us to more directly assess the speed of adjustment using an ECM model. Our approach can be interpreted as “semi-structural”: as we describe below, we instrument the sufficient statistic (the employment rate) with proxies for local demand shocks. An alternative “reduced form” approach would be to write utility as a direct function of population and all the shocks affecting the local economy.

⁵Glaeser and Gyourko (2005) argue housing supply is inelastic in cities with declining population because housing is durable. However, population decline in our sample of CZs is largely confined to small rural areas. Among the 50 largest CZs in 1950, only one (Pittsburgh) lost population in the subsequent six decades. Urban population decline is typically associated with suburbanization (away from central cities) within broader CZs.

⁶That is, it only matters for the first stage of our ECM empirical model (11) - but as we show, our instruments yield sufficient power in the first stage without accounting for this.

⁷The 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 popular such index is compiled by the American Chamber of Commerce Research Association (ACCRA), based on the prices of 59 items across 300 cities. But this has been shown to suffer from some serious deficiencies: see e.g. Koo, Phillips and Sigalla (2000); Albouy (2008); Phillips and Daly (2010).

In this section, we have presented a simple framework which underlies our empirical model. Online Appendix A shows how the same type of equation as (8) can be derived from more complex models with many traded sectors, a non-traded 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, Kolko and Saiz, 2001; Diamond, 2016), frictional labor markets (e.g. Beaudry et al., 2012; 2014a; 2014b) and heterogeneous skills (e.g. Moretti, 2011; Diamond, 2016). In Section 4.3, we consider the implications of modeling housing equity: this could theoretically yield an independent impact of house prices on migration (for a given employment rate); but the sign of this effect is ambiguous, and it seems empirically unimportant.

Next, we combine our model for short run equilibrium (with fixed population) with a simple specification for the migratory dynamics, where population responds to the gap between local utility u_r and aggregate utility u :

$$\begin{aligned} \frac{\partial l_r(t)}{\partial t} &= f(u_r(t) - u(t)) \\ &= \gamma [\tilde{a}_r(t) + n_r(t) - l_r(t)] \end{aligned} \tag{9}$$

where t denotes time. The second line substitutes (7) for $u_r(t)$, with $\tilde{a}_r(t)$ denoting a combination of area and time fixed effects. There are various models of migration (with explicit moving costs) which could explain why the migratory response is sluggish (i.e. γ less than infinite); see Kennan and Walker (2011) for a recent analysis and Online Appendix B for further discussion. As these models are typically observationally equivalent for our data, we do not seek to explain this sluggishness. Rather, our question is whether this sluggishness can help account for the persistence of local joblessness. (9) is written as a myopic model, but Online Appendix B.1 shows how an equivalent equation can be derived if agents are forward-looking and the local utility level is expected to follow a mean-reverting process.⁸ Also, (9) is agnostic about whether it is the currently employed or unemployed who are moving. As we show in Online Appendix B.2, these groups will have slightly different levels of utility; but these differences will be small, given the size of labor market flows. Consequently, the incentives to move for both groups will depend on the employment rate.

In the steady-state, “spatial arbitrage” ensures utility is equal everywhere. A positive

⁸The adjustment parameter γ would then depend on the persistence of utility differentials. Since we do not seek to offer a deep structural interpretation of the coefficients in our population growth equation, this does not affect the validity of what we do.

local demand shock initially drives up real wages and the employment rate. But population also increases, pushing real wages back down again. If the shock is small enough to have no effect on the aggregate economy, utility is unaffected in the long run, with population growth eventually ensuring higher local wages are perfectly offset by house prices. But population adjusts sluggishly, and this process is encapsulated in (9).

As we show in Online Appendix B.3, a discrete-time approximation of (9) gives:

$$l_r(t) - l_r(0) = \left[1 - \left(\frac{1 - e^{-\gamma t}}{\gamma t}\right)\right] [n_r(t) - n_r(0) + \tilde{a}_r(t) - \tilde{a}_r(0)] + (1 - e^{-\gamma t}) [n_r(0) - l_r(0) + \tilde{a}_r(0)] \quad (10)$$

assuming a constant rate of adjustment, which suggests the following empirical specification:

$$\Delta l_{rt} = \beta_0 + \beta_1 \Delta n_{rt} + \beta_2 (n_{rt-1} - l_{rt-1}) + \beta_3 \Delta \tilde{a}_{rt} + \beta_4 \tilde{a}_{rt-1} + \varepsilon_{rt} \quad (11)$$

where t denotes time periods at decadal intervals, Δ is a decadal change, $\Delta \tilde{a}_{rt}$ and \tilde{a}_{rt-1} are observed supply-side changes and levels respectively (driven by amenities or the labor supply shifter), and ε_{rt} includes any supply effects which are unobserved. Note the \tilde{a}_{rt} contain a full set of time effects reflecting changes in the aggregate level of utility in (9), and we also consider specifications which include area r fixed effects - to absorb unobserved time-invariant supply effects.

(11) is an ECM, with the change in log population Δl_{rt} responding to the (demand-driven) change in log employment Δn_{rt} and the lagged log employment rate $n_{rt-1} - l_{rt-1}$ (initial disequilibrium). If population adjusts instantaneously, β_1 and β_2 would both equal 1 (though if $\beta_1 = 1$, it would not be possible to identify β_2 - since there would be no deviations from equilibrium). β_1 and β_2 are not informative of deep structural parameters: they depend on various (unobservable) migratory frictions, as well as any expected persistence in the utility differentials. However, they do describe how local population performs in the “race” against employment - and this is what concerns us here. In (10), β_1 and β_2 are determined by a single parameter, γ . We choose not to impose this restriction as more general models would not predict it: e.g. if there is heterogeneity across workers in mobility rates, one might expect that β_1 reflects the response of the more mobile individuals and β_2 the less mobile.

This ECM can be interpreted as a generalization of the models usually estimated in the literature. Like other studies of decadal changes (e.g. Autor, Dorn and Hanson, 2013),

we control for contemporaneous shocks and identify these with suitable instruments. But this work typically excludes the disequilibrium term, implicitly assuming the local economy is in equilibrium at each observation. We claim these dynamics matter even over decadal horizons because demand shocks are very persistent. On the other hand, Blanchard and Katz (1992) do account for dynamics in their annual VAR model - but they do not control for contemporaneous shocks, which clearly matter for the decadal intervals which interest us.

3 Data

3.1 Local population and employment

We identify local labor markets r with commuting zones (CZs), collections of counties with strong commuting ties, originally developed by Tolbert and Sizer (1996) and popularized by Autor and Dorn (2013) and Autor, Dorn and Hanson (2013). We restrict our analysis to the 722 CZs of the continental US, over the period 1950-2010. 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 demographic groups. These count variables are largely based on county-level aggregates from the census (until 2010) and pooled 2009-11 American Community Survey (ACS) samples (for our 2010 cross-section), extracted from the National Historical Geographic Information System (Manson et al., 2017). The published aggregates do not cover all variables of interest, so we supplement these where required with information from census and ACS micro-data samples, taken from the Integrated Public Use Microdata Series (Ruggles et al., 2017). We describe how we construct our variables in Online Appendix D.1.

As we explain below, credible identification of (11) requires an instrument which is uncorrelated with unobserved supply shocks in the error term ε_{rt} . In keeping with much of the literature⁹, we use an industry shift-share variable b_{rt} originally proposed by Bartik (1991). This is the predicted growth of local employment (over one decade), assuming employment in each industry i grows in line with the national rate in all areas r . See Online Appendix D.2 for details on data construction. Specifically:

⁹See, for example, Blanchard and Katz (1992); Bound and Holzer (2000); Notowidigdo (2011); Beaudry et al. (2012; 2014a; 2014b).

$$b_{rt} = \sum_i \phi_{rt-1}^i \left[n_{i(-r)t} - n_{i(-r)t-1} \right] \quad (12)$$

where ϕ_{rt-1}^i is the share of employed individuals in area r at time $t - 1$ working in a 2-digit industry i , and $\left[n_{i(-r)t} - n_{i(-r)t-1} \right]$ is the change in log national employment in industry i , excluding area r .¹⁰ In Online Appendix A.2, we show how this instrument can be derived from a multi-sector version of our model.

There may be concern that the local industrial composition ϕ_{rt-1}^i lagged one decade may not be sufficiently exogenous to local supply shocks. An alternative approach is to construct the shift-shares in all years based on local industrial composition at a single historical date. We have experimented with using 1940 as our base year, but this makes little difference to our estimates (see Online Appendix G.5). This is presumably because of strong persistence in local industrial composition, which we discuss further in Section 5.4. In the same online appendix, we also study the implications of decomposing the Bartik instrument into broad industry components (agriculture/mining, manufacturing and services): again, our estimates of the population response look similar.

We instrument current employment growth Δn_{rt} using the contemporaneous Bartik shift-share b_{rt} ; and we instrument the lagged employment rate $(n_{rt-1} - l_{rt-1})$ using the lagged shift-share b_{rt-1} . The intuition for the lagged instrument is that the employment rate, at any point in time, can be written as a distributed lag of past labor demand shocks. In practice, it is sufficient to instrument using the first lag alone.

3.2 Supply controls

We control for a range of observable amenities in place of $\Delta \tilde{a}_{rt}$ and \tilde{a}_{rt-1} in (11). See Online Appendix D.3 for details on data sources. First, we include a binary indicator for presence of coastline (ocean or Great Lakes) borrowed from Rappaport and Sachs (2003). Following Rappaport (2007), we control for maximum January temperature, maximum July temperature and mean July relative humidity. We also control for log population density (measured in 1900 - to ease concerns over endogeneity) and log distance to the closest CZ, where distance is measured between population-weighted centroids. Since the impact of these amenities may vary over time (as Rappaport and Sachs, 2003, and Rappaport, 2007, argue), we interact each of them with a full set of year effects.

¹⁰This modification to standard practice, recommended by Goldsmith-Pinkham, Sorkin and Swift (2017), was originally proposed by Autor and Duggan (2003) to address concerns about endogeneity to local employment counts.

We do not control for amenities which are likely to be endogenous to current labor market conditions, such as crime and local restaurants. As we discuss in Online Appendix A.4, this means the β_1 and β_2 coefficients in equation (11) account for *all* effects of employment on local population growth, both through direct labor market effects (discussed in Section 2 above) and any indirect effects due to changes in local amenities (see Diamond, 2016).

An important contributor to local population growth is foreign migration. Of course, inflows from abroad are partly a response to demand. But as is well known, migrants are often guided in their location choice by the “amenity” of established co-patriot communities.¹¹ In the empirical literature, there is a long tradition (popularized by Altonji and Card, 1991, and Card, 2001) of proxying these preferences with historical settlement patterns - and using a Bartik-style shift-share to predict the migration shock to an area. Details of this variable’s construction are left to Online Appendix D.3.

The supply controls in (11) should also include factors that shift the labor supply curve, such as welfare allowances, state-specific supplements to Earned Income Tax Credit and minimum wages. It turns out that controlling for these has little effect on our results, whether we condition on policy levels or changes. We choose not to include them in the specifications in the main text because we do not have data for the full set of years - but Online Appendix G.2 offers estimates of the population response when they are included.

3.3 Estimates of employment rate persistence

We now offer some estimates of 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, based on all CZs and census years. 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). It is stronger among the full working-age population (row 1) than the labor force (row 4), largely due to the strong persistence of labor force participation (row 17);¹² see Yagan (2017), who also emphasizes the role of the participation margin. And it is largely driven by non-graduates (less than four years of college), especially at the larger lags (rows 5 and 6).

We claim that these results reflect persistent differences in economic opportunity for

¹¹E.g. because of job networks (Munshi, 2003) or cultural amenities (Gonzalez, 1998).

¹²Note the autocovariance of the overall employment rate can be written as the sum of the autocovariances of (i) the employment rate among the labor force and (ii) the participation rate, as well as (iii) “off-diagonal” terms, specifically the covariances between the labor force employment rate and lagged participation rate, and vice versa. The autocovariance is largely driven by (ii), especially at the larger lags.

Table 1: The autocorrelation function of the log employment rate

Employment rate variant		Lag					
		1	2	3	4	5	6
(1)	Emp rate (time-demeaned)	0.86	0.79	0.72	0.62	0.56	0.52
	<i>Sub-samples</i>						
(2)	Years 1950-80	0.87	0.81	0.72			
(3)	Years 1980-10	0.85	0.73	0.73			
(4)	Labor force	0.55	0.46	0.47	0.39	0.36	0.28
(5)	College graduate	0.37	0.25	0.16	0.08	-0.01	-0.05
(6)	Non-graduate	0.81	0.72	0.64	0.51	0.43	0.39
(7)	Male	0.79	0.71	0.68	0.57	0.51	0.25
(8)	Female	0.90	0.78	0.67	0.54	0.40	0.42
(9)	Composition-adjusted	0.83	0.74	0.67	0.58	0.47	0.39
(10)	CZ amenity controls	0.87	0.81	0.76	0.64	0.57	0.46
(11)	Within-state	0.79	0.68	0.58	0.42	0.35	0.28
(12)	Collapsed to state	0.82	0.75	0.69	0.58	0.53	0.51
	<i>Within-CZ</i>						
(13)	Unadjusted	0.33	-0.08	-0.28	-0.62	-0.48	-0.47
(14)	Bias-corrected: $\pi = 0.9$	0.79	0.66	0.58	0.40	0.35	0.31
(15)	Bias-corrected: $\pi = 0.5$	0.71	0.53	0.41	0.17	0.10	0.05
(16)	Bias-corrected: $\pi = 0$	0.69	0.51	0.38	0.13	0.05	0
(17)	Participation rate	0.89	0.82	0.74	0.64	0.57	0.60

This table summarizes autocorrelation functions of the time-demeaned log employment rate, across six decadal lags. 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 after adjusting for local demographic composition, row 10 after adjusting for observed amenities, row 11 after controlling for state fixed effects, and row 12 reports the basic ACF for state-level data. Rows 13-16 report ACFs controlling for CZ fixed effects; see Online Appendix F for further details. Finally, row 17 sets out the ACF for the time-demeaned log labor force participation rate.

individuals of fixed characteristics. There are three principal challenges to this contention, and we attempt to address these here. First, one might think the persistence is driven by geographical variation in women’s preferences for labor market participation. The ACF at smaller lags is indeed larger for women than men (see rows 7 and 8), but the difference is not large: the correlation coefficients are 0.90 and 0.79 respectively at the first lag.

Second, local variation in demographic composition might be responsible. In particular, regional skill differences are very persistent over time, and low skilled workers have lower

employment rates (Rappaport, 2012). Also, the low skilled may sort into cities with poor employment opportunities and cheap housing (Glaeser and Gyourko, 2005; Notowidigdo, 2011). But, adjusting local employment rates for demographic composition¹³ (age, education, gender and ethnicity) makes remarkably little difference to the result (see row 9). Of course, it is not possible to rule out sorting on unobservable traits (such as preferences for leisure); but the strong evidence on observables casts doubt on its importance. In Online Appendix F.3, we control for the presence of higher education institutions: this may proxy for local variation in the quality of human capital not captured by the education variable in the census, but we find it has little effect on our ACF estimates.

A third possibility is that permanent differences in amenities may compensate individuals for the persistent employment rates. We test this by purging log employment rates of our full set of supply controls; see Online Appendix F.1 for further details. But again, this makes little difference to the observed persistence (row 10). Though we believe we have controlled for the most important permanent amenities, there may be some local variation which is unobserved. But controlling for a full set of state fixed effects¹⁴ 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.

In rows 13-16, we include CZ fixed effects. Given the panel is short (only seven observations per CZ), these fixed effects are estimated with substantial error. This creates a downward bias in the ACF, as purging these effects introduces an artificial negative correlation between the employment rate and its lags. This is clear from the implausibly negative correlation coefficients in row 13. Fortunately, this bias is quantifiable. Online Appendix F.2 shows how one can derive bias-corrected estimates of the true ACF from the sample ACF, though one 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 basic ACF in row 1. And even for $\pi = 0$, the correlation is strong at the first few lags - though the fixed effects (artificially) explain much of the correlation

¹³To do this, we run logit regressions of employment on a detailed range of individual 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, separately for each census cross-section. We then predict the average employment rate in each location - assuming the local demographic composition in each location is identical to the national composition. See Online Appendix F.1 for further details.

¹⁴Some CZs straddle state boundaries, so we allocate these to the state accounting for the largest population share of the CZ.

thereafter.

We conclude that there is considerable persistence in employment rates that cannot be explained by differences in demographic composition or amenities.

4 Estimates of the population response

4.1 Baseline estimates

Panel A of Table 2 presents OLS and IV estimates of (11) using three specifications: (i) a “basic” specification which includes the standard amenity controls; (ii) a fixed effect specification, where the time-invariant component of the supply effects \tilde{a}_{rt-1} and $\Delta\tilde{a}_{rt}$ is modeled as an area fixed effect; and (iii) a first differenced specification (which also eliminates this 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. Throughout, we include the amenity controls and their interactions with year effects, though time-invariant variables are of course omitted from (ii) and (iii). We weight all observations by the lagged local population share and cluster standard errors by CZ.

The fixed effect and first differenced specifications are empirically very demanding, and other studies which analyze decadal or long term differences (such as Bound and Holzer, 2000; Autor, Dorn and Hanson, 2013; Beaudry, Green and Sand, 2014*b*) have not estimated them. Given the short time dimension (just six observations), it is difficult to empirically disentangle a “genuine” supply fixed effect (driven by local variation in the amenity or labor supply curve) from persistent joblessness 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. For this reason, we prefer the basic specification estimates.

In all three OLS specifications, the coefficient signs are what we would expect. Higher contemporaneous employment growth is associated with higher population growth; and our β_1 estimate varies little across specifications, hovering around 0.8. The OLS coefficient on the lagged employment rate is positive, implying that areas that are doing better initially tend to gain population; but this effect is much more sensitive to specification, ranging from 0.17 to 0.96.

But there is good reason to think the OLS estimates suffer from various biases. First,

Table 2: Baseline estimates of population response

PANEL A: OLS and IV						
	OLS			IV		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log \text{ emp } 16-64$	0.814 (0.012)	0.806 (0.014)	0.831 (0.012)	0.702 (0.031)	0.889 (0.052)	0.748 (0.035)
Lagged log emp rate 16-64	0.171 (0.014)	0.513 (0.031)	0.960 (0.027)	0.392 (0.056)	1.223 (0.256)	0.782 (0.165)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

PANEL B: First stage						
	$\Delta \log \text{ emp } 16-64$			Lagged log emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Current Bartik	0.972 (0.074)	0.930 (0.079)	0.756 (0.071)	0.041 (0.040)	-0.111 (0.035)	-0.020 (0.028)
Lagged Bartik	0.094 (0.059)	-0.024 (0.059)	-0.118 (0.072)	0.453 (0.046)	0.131 (0.035)	0.150 (0.022)
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 equation (11), as well as first stage estimates (Panel B), across 722 CZ over 1950-2010. The dependent variable is the decadal change in log population, and the endogenous variables are the change in log employment and the lagged log employment rate, for individuals aged 16-64. These are instrumented with the current and lagged Bartik shift-shares. Throughout, we control for year effects, amenity measures (interacted with all the year effects) and the migrant shift-share. Columns 2 and 5 also control for CZ fixed effects, and columns 3 and 6 report a first differenced specification. Robust standard errors, clustered by CZ, are reported in parentheses. Observations are weighted by the lagged local population share.

population and employment growth are jointly determined¹⁵, so the coefficient on employment growth cannot be interpreted as causal. Specifically, unobserved supply-side shocks to population (due to e.g. local amenities) will affect local job creation, and this should bias the OLS estimate of β_1 upwards. Furthermore, if these unobserved supply shocks are persistent, OLS estimates of β_2 may be biased downwards. For example, an improvement in local amenities should affect local population growth positively but the employment rate negatively. To the extent that these amenity effects are persistent, some of these biases may be addressed somewhat in the fixed effect specification. But, the inclusion of fixed effects may introduce a “Nickell bias” (Nickell, 1981): demeaning creates 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 the variables of interest using current and lagged Bartik shift-shares, as described in Section 3 above. The identifying assumption is that the lagged local industry shares underlying the Bartik shift-share (i.e. the ϕ_{rt-1}^i in (12)) are exogenous to omitted amenities and labor supply shifters: see Goldsmith-Pinkham, Sorkin and Swift (2017) for a useful discussion. And when we control for CZ fixed effects or take first differences, it is the interaction of these local shares with aggregate-level industry trends that our instrument exploits.¹⁶

We would expect the current Bartik shift-share to be most strongly correlated with current employment growth and the lagged shift-share 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 power¹⁷ and sufficient independent variation despite substantial serial correlation in the Bartik shift-share (see Section 5.4 below). Contemporaneous employment growth is *only* responsive to the current Bartik shift-share, with a coefficient close to 1. And the lagged employment rate is *only* responsive to the lagged instrument, with a coefficient of about 0.5 in the basic specification. The smaller effect of the latter is indicative of a sizable population response, consistent with our second stage results. When we introduce fixed effects or estimate in first differences¹⁸, the patterns of correlation between the endogenous variables and the instruments become more

¹⁵We estimate the employment response to population in Section 5 below.

¹⁶In principle, the instrument is also exploiting local changes in industry shares over time: it is based on local shares at a single lag; see (12). But as Online Appendix G.5 shows, fixing the shares at 1940 levels makes little difference to the estimates.

¹⁷The Cragg-Donald test statistics for a null of weak instruments are 108, 37 and 58 for the basic, fixed effect and first differenced specifications respectively. These are well in excess of the critical value of 7.03 for two endogenous variables and two instruments (Stock and Yogo, 2005), for a maximal IV size of 10 percent.

¹⁸The instruments in the first differenced specification are the first differenced Bartik shift-shares, both current and lagged.

complicated - though the instruments remain powerful.

The IV estimates are reported in columns 4-6 of Table 2; and overall, the model performs well in all specifications. At least in the basic specification, using IV shifts the estimates in the expected direction compared to OLS. The lagged error-correction term, a novel feature of our specification, is strongly significant in all specifications in Table 2, taking a value of 0.39 in our (preferred) basic specification. We show in Online Appendix G.1 that omitting this term yields a larger coefficient on contemporaneous local employment growth: 0.85 compared to the 0.70 in the basic specification of column 4. At least in relation to a value of 1 (commensurate with a complete population response), this is an important difference. This is indicative of serial correlation in the Bartik shift-share: we return to this in Section 5.4.

Our IV estimates of β_1 and β_2 are closer to 1 under fixed effects and first differences, though they look more similar to the basic specification in various robustness exercises in Online Appendix G.¹⁹ More broadly, this appendix shows our IV results are robust to different choices about weighting, controls (for amenities, state welfare policies, the presence of local colleges, and predicted shifts in local industry rents), instruments and outliers. In Online Appendix I, we show the population response is largely driven by variation in migratory inflows rather than outflows. This is consistent with earlier work by Coen-Pirani (2010) and Monras (2015a).

Broadly speaking, Table 2 reveals a robust relationship between population growth on the one hand and employment growth and the lagged employment rate on the other.²⁰ One way to summarize the results is to consider the response to a fixed change in employment, from an initial equilibrium. Based on the basic IV estimates (column 4), a 10 percent contraction of local employment pushes population down by 7.0 percent in the first decade - a large change, but not sufficient to remove the impact of the shock. As the employment rate is 3.0 percent lower after 10 years, population continues to decline by 1.2 percent in the second decade (based on our β_2 estimate). So after 20 years, the employment rate would be 1.8 percent below the pre-shock level.

¹⁹Specifically, they are closer to the basic specification when we do not weight the observations by population (Online Appendix G.1) and when we disaggregate the Bartik instruments into broad industry components (Online Appendix G.5).

²⁰Our ECM exploits our claim that the employment rate can be used as a sufficient statistic for economic opportunity. But as we noted earlier, one might also use the real wage. Online Appendix J shows one can obtain sensible estimates of the population response to real wages, but these are sensitive to assumptions about the importance of housing costs in consumer price indices. If both the real wage and employment rate are included, then the latter appears more important in practice. We take this as evidence that our approach is sensible, but that one could reach similar conclusions using real wages.

Table 3: Heterogeneity in IV population responses

	1950-1980	1980-2010	Lab force	Coll grad	Non grad	16-24s	25-44s	45-64s
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Basic specification</i>								
Δ log emp	0.811 (0.038)	0.393 (0.055)	0.880 (0.018)	0.913 (0.041)	0.673 (0.036)	0.613 (0.033)	0.788 (0.037)	0.660 (0.043)
Lagged log emp rate	0.247 (0.076)	0.573 (0.095)	1.371 (0.336)	1.037 (0.269)	0.456 (0.069)	0.431 (0.043)	0.506 (0.084)	0.356 (0.092)
<i>CZ fixed effects</i>								
Δ log emp	0.918 (0.042)	0.428 (0.065)	1.041 (0.114)	0.894 (0.048)	0.855 (0.071)	0.768 (0.058)	0.905 (0.039)	0.881 (0.097)
Lagged log emp rate	0.757 (0.236)	0.615 (0.117)	4.539 (3.429)	0.731 (0.125)	1.660 (0.460)	0.923 (0.168)	2.028 (0.687)	1.371 (0.571)
<i>First differences</i>								
Δ log emp	0.885 (0.048)	0.149 (0.152)	0.883 (0.022)	0.782 (0.116)	0.709 (0.034)	0.619 (0.036)	0.821 (0.027)	0.760 (0.051)
Lagged log emp rate	0.500 (0.461)	0.214 (0.232)	1.265 (0.288)	1.176 (0.335)	0.953 (0.195)	0.582 (0.132)	1.388 (0.223)	1.202 (0.258)
Obs (basic, FE)	2,166	2,166	4,332	4,331	4,332	4,332	4,332	4,332

Each column reports IV estimates of β_1 and β_2 in the population response equation (11) for a different subsample. Columns 1 and 2 report estimates for 1950-1980 and 1980-2010 respectively. In column 3, population is replaced by labor force, and the employment rate is measured as a share of labor force participants only (i.e. excluding the inactive). Columns 4 and 5 split the sample by education, and columns 6-8 by age. In columns 4-8, all variables and instruments are constructed using group-specific data. For other columns, variables and instruments are based on individuals aged 16-64. Observation counts for the basic and fixed effect specifications are given in the final row. The first differenced sample is 772 smaller in each case. Column 4 is missing one observation, because in one largely rural CZ (centered around Mecosta County MI), there were no working-age employed graduates in the 1950 micro-data extract. See notes under Table 2 for further details on empirical specification and right hand side controls. Robust standard errors, clustered by CZ, are reported in parentheses. Observations are weighted by the lagged local population share.

Overall, the evidence suggests our model for population growth performs well and is a useful framework for studying the response to employment shocks.

4.2 Heterogeneity in responses

The estimates above describe aggregate population responses, across six decades of data and all working-age individuals. Table 3 reports heterogeneity in the IV estimates by time, economic activity, education and age. We leave the corresponding OLS and first stage estimates to Online Appendix H. Also, we do not consider origin country here - but using the same data and empirical model, Amior (2017*a*) shows that foreign migrants make a disproportionate contribution to local population adjustment.²¹

The first two columns of Table 3 report estimates for before and after 1980. Gross migration rates have declined since 1980 (see Molloy, Smith and Wozniak, 2011, 2017; Kaplan and Schulhofer-Wohl, 2017), and there is concern that population may have become less responsive to local shocks (Partridge et al., 2012; Beyer and Smets, 2015; Dao, Furceri and Loungani, 2017). The fixed effect and first differenced specifications do suggest the population response is weaker after 1980, though these specifications are especially demanding given a time dimension of just four observations. Having said that, the evidence from the basic specification is mixed: while the IV coefficient on employment growth is smaller after 1980, the coefficient on the lagged employment rate is larger. It is also worth noting that the persistence of the local employment rate changes little after 1980: see Table 1.

Next, we turn to economic activity. Our model in Section 2 abstracts from the participation decision: it does not distinguish between working-age population and labor force. The third column of 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 from 1.²² The implication is that any sluggishness in the population response to initial local employment rate differentials is entirely manifested on the participation margin. The importance of the participation margin tallies with findings from Black, Daniel and Sanders (2002), Autor and Duggan (2003), Rappaport (2012) and Autor, Dorn and Hanson (2013), who identify large declines in participation (and rising take-up of disability benefits) in cities subject to adverse shocks. This is also consistent with the finding that the employment rate is less persistent among labor force participants than the broader population: see Table 1.

²¹This is consistent with Borjas (2001) and Cadena and Kovak (2016).

²²This is in line with the results of Beaudry, Green and Sand (2014*b*), who restrict their sample to the labor force.

In columns 4 and 5, we estimate population responses separately for college graduates and non-graduates. As the theory would suggest (see Online Appendix A.6), we use group-specific stocks for all relevant variables: population growth, employment growth and the lagged employment rate. We construct the industry shift-share instruments using group-specific employment counts (see Online Appendix D.2), and we also construct our migrant shift-share controls separately for each group (Online Appendix D.3). In line with other work (e.g. Bound and Holzer, 2000; Wozniak, 2010; Notowidigdo, 2011), we estimate larger population responses among better educated workers, at least in our (preferred) basic specification.²³ This is also consistent with our finding in Table 1 that their employment rates are less persistent. In our basic IV results, β_1 is 0.91 for graduates and 0.67 for non-graduates. And the elasticity β_2 to the lagged employment rate is 1.04 for graduates and 0.46 for non-graduates; though β_2 is larger among non-graduate workers in the IV fixed effect specification.

Finally, in columns 6-8, we study heterogeneous population responses by age. The equation works well for all age groups, though the 25-44s respond somewhat more strongly than older or younger individuals. Despite this, we show in Online Appendix H that the labor force responses are similarly large across age groups. This is consistent with the view that, following adverse shocks, older workers disproportionately drop out of the labor force rather than relocate.

The coefficients on the lagged employment rate are more varied among the fixed effect and first differenced specifications in Table 3. But these specifications are more demanding empirically, and this is reflected in the larger standard errors.

4.3 Implications of housing wealth

Our baseline model implicitly assumes all workers are renters, so housing costs only affect welfare through the real consumption wage - or, by the sufficient statistic result, through the employment rate. But if some residents are owner-occupiers, our sufficient statistic result may be violated: changes in house prices may now affect population even after controlling for the employment rate. But the sign of this effect is theoretically ambiguous; and empirically, it does not seem very large.

²³It is often argued that workers with less education face higher migration costs (Bound and Holzer, 2000; Wozniak, 2010). Notowidigdo (2011) claims they are relatively sheltered from local demand shocks because of declining housing costs and transfer payments. Amior (2017*b*) argues the major obstacle to their mobility is the meager wage rents accruing to their job matches. It should be stressed however that our estimates in Table 1 focus on education-specific population changes, so they conflate both mobility and local cohort effects.

The ambiguity arises from two offsetting effects. First, housing wealth may exert a direct effect on utility - conditional on the local employment rate. If housing wealth is transferable geographically, a fall in local house prices means that someone considering moving to another area can now only afford a lower level of housing services in potential destinations. This discourages migration from areas with declining relative house prices.²⁴ The effect may be amplified if declining prices push households into negative equity (see e.g. Ferreira, Gyourko and Tracy 2010; Foote 2016). On this basis, conditional on the local employment rate, we would expect population to be decreasing in local house prices - relative to elsewhere. Thus, the employment rate would no longer serve as a sufficient statistic for utility.

But housing wealth may also exert an indirect effect through labor supply. Local house price appreciation would expand housing wealth, and this would shift the labor supply curve downwards if leisure is a normal good. A low employment rate may then not necessarily reflect low real wages: it could also mean that housing equity has grown. House prices would then enter utility positively, conditional on the employment rate - thus offsetting the direct wealth effect.

If the sufficient statistic result is indeed violated, this would yield an ECM for population adjustment in which the change in log house prices and its lagged level are included as right hand side controls (following the reasoning in Section 2). We present IV estimates of such an equation in Table 4. We use a residualized index of local prices of owner-occupied housing, purged of variation in observed housing characteristics, based on census and ACS micro-data (see Online Appendix D.4 for details). The price variables are clearly endogenous, and we instrument them using Saiz's (2010) estimates of the elasticity of housing supply and its interactions with current and lagged Bartik shift-shares. Intuitively, labor demand shocks should have a smaller impact on house prices in areas where housing is supplied more elastically. At least in the basic specification, this prediction is supported by the first stage estimates in Panel B. Our sample is restricted to the period since 1960 (since the 1950 census does not report house prices) and the 248 CZs for which we have elasticity estimates from Saiz (see Online Appendix D.4).

Columns 1-3 of Panel A reproduce our main specification (11) for the restricted sample, and the results are similar to Table 2. Columns 4-6 then include the change in residualized log house prices and the lagged log price. The associated coefficients are mostly insignificantly

²⁴Intuitively, as Sinai and Souleles (2005) have argued, home-ownership can serve as a form of insurance against fluctuations in rents: housing assets effectively pay dividends which offset the implicit rental costs. However, households do not benefit from this insurance if they move to another location whose prices are uncorrelated with their initial place of residence (Sinai and Souleles, 2013).

Table 4: IV estimates of population response controlling for house prices

PANEL A: IV estimates						
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log \text{ emp}$	0.646 (0.052)	0.705 (0.075)	0.682 (0.054)	0.612 (0.085)	0.902 (0.151)	0.855 (0.151)
Lagged log emp rate	0.398 (0.100)	1.507 (0.498)	0.873 (0.244)	0.382 (0.232)	1.250 (0.484)	0.507 (0.388)
$\Delta \log \text{ price}$				0.027 (0.077)	-0.251 (0.142)	-0.068 (0.068)
Lagged log price				0.010 (0.041)	-0.054 (0.106)	0.161 (0.223)
Observations	1,240	1,240	992	1,240	1,240	992

PANEL B: First stage for house prices						
	$\Delta \log \text{ price}$			Lagged log price		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Current Bartik	1.322 (0.288)	0.820 (0.444)	0.483 (0.680)	-0.146 (0.346)	-1.385 (0.376)	-0.777 (0.301)
Lagged Bartik	-0.247 (0.218)	-0.018 (0.343)	0.134 (0.683)	2.070 (0.470)	0.905 (0.241)	0.482 (0.200)
Current Bartik * elasticity	-0.103 (0.060)	0.086 (0.085)	0.321 (0.197)	0.050 (0.113)	0.249 (0.109)	0.069 (0.075)
Lagged Bartik * elasticity	-0.080 (0.050)	-0.092 (0.066)	-0.242 (0.116)	-0.328 (0.111)	-0.132 (0.040)	-0.043 (0.044)
Elasticity	0.005 (0.008)			-0.039 (0.020)		
Observations	1,240	1,240	992	1,240	1,240	992

Panel A reports IV estimates of the population response equation (11), though in columns 4-6, we also include the change in log (residualized) house prices and the lagged house price on the right hand side. See Online Appendix D.4 for details on the residualization process. Panel B reports the first stage for changes in house prices and for lagged prices. The sample is restricted to the five (decadal) periods since 1960 (since the 1950 census does not report house prices) and the 248 CZs for which we have housing supply elasticity estimates from Saiz (2010). See notes under Table 2 for details on the right-hand side controls. Errors are clustered by CZ, and robust standard errors are reported in parentheses. Each observation is weighted by the lagged local population share.

different from zero, and the effects of employment are similar to columns 1-3. So it does not appear that accounting for housing wealth alters our substantive conclusions.

5 Dynamic response to demand shocks

5.1 Estimates of employment response

The estimates above allow us to predict the population response, conditional on local employment. But, this is not the same as the response to a labor demand shock, because employment itself is endogenous. If the adjustment of labor demand is sluggish (like population), we show in Online Appendix C that we can derive a similar ECM equation for employment growth:

$$\Delta n_{rt} = \alpha_0 + \alpha_1 \Delta l_{rt} + \alpha_2 (n_{rt-1} - l_{rt-1}) + \alpha_3 b_{rt} + d_t + \omega_{rt} \quad (13)$$

where b_{rt} is the Bartik shift-share, the d_t are time effects, and ω_{rt} is an unobserved component. This is similar to the population response equation, but with population and employment reversed. In the simple model above, population affects employment only through wage changes. But in practice, an expanding population will also raise demand for non-traded goods: see, for example, Moretti (2010) and Howard (2017), or the model extension in Online Appendix A.1. And if the demand response is sluggish, employment will also depend on the lagged employment rate - which accounts for deviations from equilibrium (see Online Appendix C). This equation considers employment as a whole, though a disaggregation into number of establishments and workers per establishment might also be of interest: see Online Appendix K for empirical estimates.

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 b_{rt} in the regression. But, it is harder to identify a convincing instrument for population growth which is exogenous to changes in labor demand. Our strategy is to use the maximum January temperature (also interacted with a time trend) as our instrument, as Americans have increasingly been attracted to cities with mild winters (Rappaport, 2007). We control for all the other usual amenity controls on the right hand side, including maximum July temperature and mean July humidity on the right-hand side: hot summers may conceivably be linked to growing labor demand, if the expansion of air conditioning increased labor productivity (Oi, 1996). Beaudry et al. (2014a; 2014b) also use climate instruments as well

as the migrant shift-share to identify supply shocks, though the latter is a weak instrument in our data.²⁵

The results are reported in Table 5. Across all specifications, the OLS estimate of α_1 (the response to population) is equal to 1. The IV elasticity in the basic specification is smaller, at 0.79.²⁶ This difference is to be expected: we know from Table 2 that population responds positively to employment.

Next consider α_2 , the coefficient on $n_{rt-1} - l_{rt-1}$. In our basic IV specification, a positive 10 percent deviation in the initial employment rate leads to a 2.1 percent decrease in subsequent employment growth. Again, this seems sensible: tighter labor markets (with higher wages) should discourage job creation.

The fixed effect and first differenced IV estimators have little power, but the identification (purely through the interaction between January temperature and the time trend) is clearly demanding. Given this, we rely on the basic IV estimates (column 4) in the analysis which follows.

5.2 Impulse response

Our estimates of the population and employment equations (11, 13) allow us to simulate the dynamic response to a local labor demand shock. Specifically, we rely on the following simultaneous equation model:

$$\Delta n_{rt} = \alpha_0 + \alpha_1 \Delta l_{rt} + \alpha_2 (n_{rt-1} - l_{rt-1}) + \Delta z_{rt}^d \quad (14)$$

$$\Delta l_{rt} = \beta_0 + \beta_1 \Delta n_{rt} + \beta_2 (n_{rt-1} - l_{rt-1}) \quad (15)$$

where Δz_{rt}^d represents the demand shock. From these two equations, one can derive a model for the evolution of x_{rt} , the deviation of the employment rate from its steady-state value:

$$\begin{aligned} x_{rt} &= \theta_1 x_{rt-1} + \theta_2 \Delta z_{rt}^d \\ &= \left[1 - \frac{(1 - \alpha_1) \beta_2 - (1 - \beta_1) \alpha_2}{1 - \alpha_1 \beta_1} \right] x_{rt-1} + \left[\frac{1 - \beta_1}{1 - \alpha_1 \beta_1} \right] \Delta z_{rt}^d \end{aligned} \quad (16)$$

²⁵This may be suggestive of geographical displacement of previous residents by new migrants: see Amior (2017a) for further analysis using the same data.

²⁶To compare, Beaudry et al. (2014a; 2014b) estimate an IV elasticity of employment with respect to the labor force of close to 1. The difference is largely explained by the choice of right hand side variable: we estimate a coefficient of 0.93 when we substitute labor force for population in equation (13).

Table 5: Estimates of employment response

PANEL A: OLS and IV						
	OLS			IV		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log \text{ pop } 16-64$	1.027 (0.011)	1.023 (0.015)	1.032 (0.014)	0.788 (0.052)	-0.297 (0.763)	3.319 (4.002)
Lagged log emp rate 16-64	-0.122 (0.012)	-0.646 (0.044)	-1.172 (0.038)	-0.207 (0.056)	0.176 (0.587)	-1.586 (1.676)
Current Bartik	0.177 (0.024)	0.111 (0.035)	0.160 (0.023)	0.425 (0.055)	1.155 (0.621)	-1.092 (2.209)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

PANEL B: First stage						
	$\Delta \log \text{ pop } 16-64$			Lagged log emp rate 16-64		
	Basic	FE	FD	Basic	FE	FD
	(1)	(2)	(3)	(4)	(5)	(6)
Max temp January	0.359 (0.082)			-0.005 (0.056)		
Max temp January * time	-0.005 (0.014)	-0.008 (0.015)	-0.025 (0.019)	-0.043 (0.015)	-0.041 (0.013)	-0.045 (0.013)
Lagged Bartik	0.249 (0.056)	0.152 (0.051)	0.037 (0.061)	0.452 (0.044)	0.136 (0.035)	0.158 (0.022)
Current Bartik	0.697 (0.064)	0.692 (0.064)	0.549 (0.055)	0.044 (0.039)	-0.107 (0.034)	-0.009 (0.027)
Observations	4,332	4,332	3,610	4,332	4,332	3,610

This table reports estimates of α_1 and α_2 in the employment response equation (13). Methods and sample are the same as in Table 2. The dependent variable is now the change in log employment, and the change in log population and current Bartik shift-share are regressors. As in Table 2, we instrument the lagged employment rate with the lagged Bartik shift-share. And we instrument local population growth with the maximum January temperature and its interaction with a time trend. Otherwise, we use the same controls as in Table 2. Robust standard errors, clustered by CZ, are reported in parentheses. Observations are weighted by the lagged local population share.

Table 6: Sensitivity of computed persistence and amplification parameters

		AR(1) persistence θ_1				Shock amplification θ_2			
		Value of α_2 :				Value of α_2 :			
		0	-0.207	-0.6	-0.9	0	-0.207	-0.6	-0.9
	0	0.608	0.546	0.429	0.340	0.298	0.298	0.298	0.298
Value	0.3	0.653	0.574	0.426	0.313	0.377	0.377	0.377	0.377
of α_1 :	0.788	0.814	0.676	0.414	0.214	0.667	0.667	0.667	0.667
	0.9	0.894	0.726	0.408	0.165	0.809	0.809	0.809	0.809

This table computes the AR(1) persistence θ_1 and the shock amplification θ_2 in equation (16), for different values of α_1 and α_2 . In each instance, we set β_1 and β_2 to 0.702 and 0.392 respectively, our estimates from the basic specification of the population equation (11) in Table 2. Notice the $\alpha_1 = 0.788$ and $\alpha_2 = -0.207$ cases correspond to our basic IV estimates of the employment equation (13) in column 4 of Table 5.

The employment rate follows an AR(1) process. We denote the persistence by θ_1 : this is dampened by larger β_1 and β_2 as well as larger (more negative) α_2 , though the sign of α_1 's effect depends on the other parameters. The shock's initial amplification, θ_2 , is increasing in α_1 but moderated by β_1 .

Our preferred estimates of the α s and β s²⁷ place the AR(1) persistence θ_1 at 0.68 and the shock amplification θ_2 at 0.67. As noted above, we are less confident about our identification of the employment equation - and α_1 specifically. In Table 6, we study the sensitivity of the persistence and amplification to α_1 and α_2 , given our preferred estimates of β_1 and β_2 . For comparison, we have included our estimates of α_1 and α_2 in the table: 0.788 and -0.207 respectively. The amplification is invariant to α_2 , though does vary markedly with α_1 . In contrast, the persistence is (reassuringly) not so sensitive to α_1 and depends much more on α_2 (whose estimate we trust more).

Significantly, turning the employment response off entirely (setting both α parameters to zero) slightly *reduces* the AR(1) persistence - from our estimate of 0.68 to 0.61. This suggests that population movements (rather than employment) account for the entire adjustment to local demand shocks, consistent with Blanchard and Katz (1992) and Hornbeck (2012). We do estimate a significant (equilibrating) response of employment to an initially slack labor market, i.e. α_2 ; but this effect is undone by the employment response to population, α_1 .

For example, consider a one-off permanent local demand shock: an innovation of 0.1

²⁷From our basic IV specifications: $\alpha_1 = 0.788$ and $\alpha_2 = -0.207$ (column 4, Panel A, Table 5), and $\beta_1 = 0.702$ and $\beta_2 = 0.392$ (column 4, Panel A, Table 2).

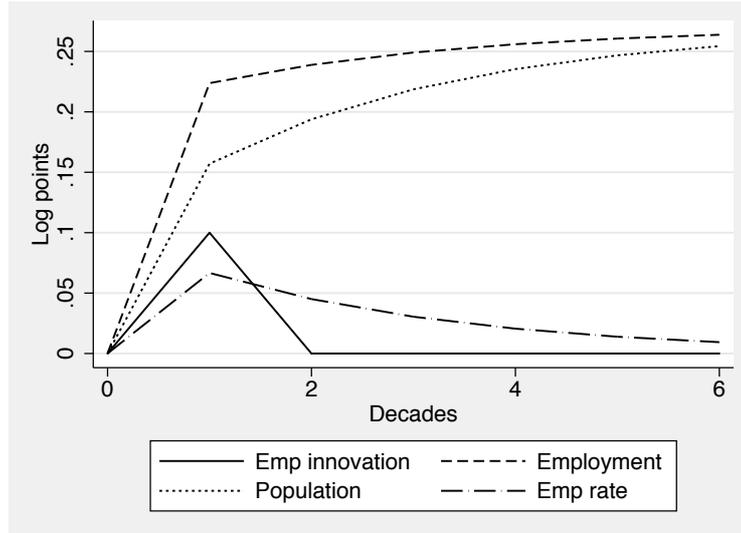


Figure 3: Impulse response function

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

log points in z_{rt}^d in area r at $t = 1$, from an initial steady-state. Figure 3 shows the impulse response, based on our preferred α and β estimates. The large employment response (to population growth) greatly amplifies the impact on local stocks, with employment and population coming to rest about 0.25 log points above their original level. In terms of the dynamics, notice employment increases somewhat following the initial shock, while population converges towards it - confirming the burden of adjustment falls on population.

Figure 3 also traces the employment rate. Population adjustment is somewhat sluggish, so the effect persists beyond one decade. Having said that, the model alone certainly cannot match the magnitude of persistence in the data, especially after the first lag or two. While equilibrium is largely restored in the model three or four decades after the shock (in Figure 3), the ACF of the local employment rate only reaches 0.5 by the sixth decadal lag (see Table 1).

5.3 Matching model to data

What explains this excess persistence in the data? In Section 3.3, we found little evidence that demographic composition or supply-side factors play a substantial role. We argue instead that persistent shocks to local demand may be responsible - as Figure 2 in the introduction suggests.

The key idea comes from a well-known feature of ECMs: long run trends can cause a permanent deviation from equilibrium. To see this, suppose the demand shock Δz_{rt}^d in (14) is fixed at some constant Δz_r^d : every period, an area experiences the same increase or decrease in labor demand, so there is an area-specific trend. In the steady-state, the local employment rate deviation x_{rt} must rest at some level x_r^* . Imposing $x_{rt} = x_r^*$ in equation (16) gives:

$$x_r^* = \frac{1 - \beta_1}{(1 - \alpha_1)\beta_2 - (1 - \beta_1)\alpha_2} \Delta z_r^d \quad (17)$$

so that areas facing a negative trend in local demand will have permanently lower employment rates²⁸ (unless $\beta_1 = 1$, i.e. immediate adjustment).

Of course, a permanent shock is an extreme case. More generally, suppose the demand shock Δz_{rt}^d follows an AR(1) process:

$$\Delta z_{rt}^d = \lambda \Delta z_{rt-1}^d + \varsigma_{rt} \quad (18)$$

where ς_{rt} is a white noise shock. We can then ask: how much persistence in Δz_{rt}^d (i.e. what value of λ) is required to match the observed persistence in joblessness? Substituting equation (18) for Δz_{rt}^d in (16) gives an AR(2) expression for x_{rt} :

$$(1 - \lambda L)(1 - \theta_1 L)x_{rt} = \theta_2 \varsigma_{rt} \quad (19)$$

where L is the lag operator, θ_1 is the model-generated persistence in x_{rt} , and θ_2 is the initial amplification of the shock. Notice in (19) that θ_1 and λ are not separately identifiable using data on the employment rate x_{rt} alone. But, we can calibrate θ_1 using our IV estimates of the α and β parameters. And we can then see what λ value is required to match the data.

We plot simulated ACFs of x_{rt} for different λ values in Figure 4, together with the ACF of the time-demeaned log employment rate, purged of observable amenity effects (the 10th row in Table 1). The model with $\lambda = 0$ (i.e. a white noise demand shock) accounts for most of the observed persistence at the first decadal lag, but its explanatory power declines substantially thereafter. To match the autocorrelation at the first lag, a λ of approximately 0.5 is required; and the sixth lag requires a λ of 0.8. A higher order process for Δz_{rt}^d would likely achieve a better fit.

²⁸For our benchmark α and β values, the coefficient on Δz_r^d is 2.06. So a 1 percent permanent deviation in demand growth yields a 2 percent employment rate deviation.

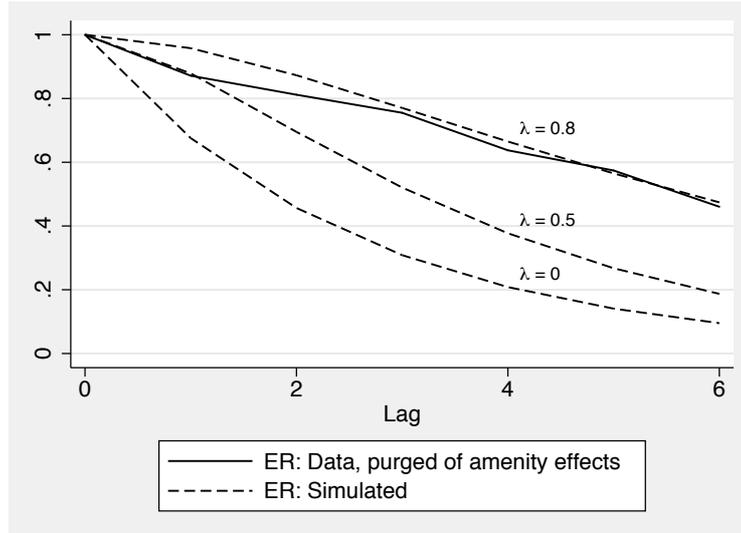


Figure 4: ACFs of employment rate: data and simulated for different λ

This figure illustrates the observed and simulated persistence of the log employment rate. The solid 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 (19), given our preferred estimates of the α and β parameters.

5.4 The plausibility of persistence

We have argued persistent joblessness is explained by persistent demand shocks, combined with some sluggishness in the population response. This section considers the plausibility of our claims. First, on the demand side, are the values of λ discussed above realistic? We do not observe all demand shocks directly, but the Bartik shift-share has an AR(1) persistence of 0.69, which falls between the 0.5 and 0.8 bounds required to match the employment rate ACF. The shift-share's persistence derives from secular declines in agriculture (Michaels, Rauch and Redding, 2012) and (since 1960) manufacturing (see Online Appendix D.2), combined with stickiness in local industrial composition.²⁹ In fact, changes to local composition do nothing to dampen the persistence: a shift-share with industry composition fixed at 1940

²⁹Individual industries are likely to have different levels of persistence, e.g. oil and gas are relatively volatile. But it is ultimately the persistence in the *overall* demand for labor that matters here, and the Bartik shock is designed to measure this. While much of the literature has relied on Bartik shocks, some studies have focused on shifts in more specific industries. For example, Autor, Dorn and Hanson (2013) study local shocks arising from industries exposed to Chinese import competition. Their work covers two time periods: 1990-2000 and 2000-2007. A regression of the local change in Chinese import exposure in the latter period on its lag, across 722 CZs, yields a coefficient of 1.2. An instructive comparison can also be drawn with Foster, Haltiwanger and Syverson (2008), who estimate a firm-level AR(1) persistence in revenue productivity of 0.3 over five year horizons (see Table 3 in that paper). This is clearly much weaker than the persistence in local demand we require. But, their estimates are picking up the large idiosyncrasies between individual firms, whereas we are interested in secular local-level trends.

(see Section 3.1) has an AR(1) parameter of 0.72, only slightly larger than when we do not fix the composition. Any persistence in the Bartik shift-shares is likely to be amplified by large local spillovers: see e.g. Greenstone, Hornbeck and Moretti (2010), Kline and Moretti (2014) and Gathmann, Helm and Schönberg (2016). The flip-side of declining manufacturing is growing demand in ideas-producing regions: see Moretti (2004), Glaeser (2005), Glaeser and Ponzetto (2007) and Moretti (2012). Related to this, Glaeser, Kerr and Kerr (2015) emphasize the role of persistent entrepreneurial clusters in driving urban growth. Dix-Carneiro and Kovak (2017) also argue that sluggish capital adjustment causes local labor demand shocks to persist.

Second, one might wonder why population does not respond more quickly - especially if people realize that demand shocks are so persistent. Although we identify a sizable population response, it is slower than much of the earlier literature suggests. Blanchard and Katz (1992), the seminal study in this field, find the effect of a state-level employment shock on the employment rate disappears within seven years.³⁰ Time horizons in the empirical model may help explain the difference in our results: without access to the data now available, Blanchard and Katz were constrained to estimating a state-level VAR with just two annual lags (in local employment growth, employment rate and participation rate), based on Current Population Survey waves between 1978 and 1990. But, Obstfeld and Peri (1998) and Jordà (2005) argue that any misspecification errors in a short run VAR may be compounded when projecting over longer horizons. For example, if the most mobile workers are the quickest to respond to an employment shock, projecting the initial responses forward may over-state the long run effect.

6 Conclusion

Joblessness is very persistent across US commuting zones. We claim this persistence cannot be explained by differences in demographic composition or permanent amenities. These disparities persist despite a strong migratory response, and we attribute this to persistent shocks to labor demand.

We also make two methodological contributions. First, using a simple Rosen-Roback model with the plausible modification that labor supply is not completely inelastic, we show how the employment rate can be used as a sufficient statistic for local economic opportunity.

³⁰Beyer and Smets (2015) and Dao, Furceri and Loungani (2017) have found comparable results using an updated dataset. See Decressin and Fatás (1995), Jimeno and Bentolila (1998) and Obstfeld and Peri (1998) for European estimates of this model.

This has advantages over the more conventional real consumer wage. Second, combining this model of local equilibrium with a dynamic migration equation yields an error correction mechanism, in which population growth is influenced by employment growth (driven by local demand shocks) and the lagged local employment rate (which measures the level of disequilibrium). Existing studies in the urban literature typically do not include the disequilibrium term, but we argue it is crucial for understanding the long-run persistence in joblessness.

We estimate our ECM model using current and lagged Bartik shift-shares as instruments, and we show the model performs well empirically. The migratory response is large, but not as large as commonly thought: adjustment to shocks is incomplete within a decade. Still, the persistence generated by our model is much weaker than what we observe in the data. We argue that serial correlation in local demand growth can account for the difference. This gives rise to a “race” between employment and population. Population does respond strongly to a decline in demand; but further demand contractions are likely to follow, so population never catches up - and the employment rate only returns very slowly to its original level. We have argued the persistence in local demand growth may be driven by secular changes in the industrial structure of employment, coupled with stickiness in local industrial composition - but the exact causes of this stickiness should be the subject of future research.

References

- Albouy, David.** 2008. “Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas.” NBER Working Paper No. 14472.
- Altonji, Joseph G., and David Card.** 1991. “The Effects of Immigration on the Labor Market Outcomes of Less-Skilled Natives.” In *Immigration, Trade, and the Labor Market*, ed. Richard B. Freeman John M. Abowd, 201–234. Chicago: University of Chicago Press.
- Amior, Michael.** 2017*a*. “The Contribution of Foreign Migration to Local Labor Market Adjustment.” <http://sites.google.com/site/michaelamior>.
- Amior, Michael.** 2017*b*. “Education and Geographical Mobility: The Role of Wage Rents.” <http://sites.google.com/site/michaelamior>.
- Autor, David H., and David Dorn.** 2013. “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market.” *American Economic Review*, 103(5): 1553–1597.
- Autor, David H., and Mark G. Duggan.** 2003. “The Rise in the Disability Rolls and the Decline in Unemployment.” *Quarterly Journal of Economics*, 118(1): 157–206.

- Autor, David H., David Dorn, and Gordon H. Hanson.** 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review*, 103(6): 2121–2168.
- Bartik, Timothy J.** 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand.** 2012. “Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory.” *Econometrica*, 80(3): 1063–1104.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand.** 2014a. “In Search of Labor Demand.” NBER Working Paper No. 20568.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand.** 2014b. “Spatial Equilibrium with Unemployment and Wage Bargaining: Theory and Estimation.” *Journal of Urban Economics*, 79: 2–19.
- Beyer, Robert C.M., and Frank Smets.** 2015. “Labour Market Adjustments and Migration in Europe and the United States: How Different?” *Economic Policy*, 30(84): 643–682.
- Black, Dan, Kermit Daniel, and Seth Sanders.** 2002. “The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust.” *American Economic Review*, 92(1): 27–50.
- Blanchard, Olivier J., and Lawrence F. Katz.** 1992. “Regional Evolutions.” *Brookings Papers on Economic Activity*, 23(1): 1–76.
- Blanchflower, David G., and Andrew J. Oswald.** 1994. *The Wage Curve*. Cambridge: MIT Press.
- Borjas, George J.** 2001. “Does Immigration Grease the Wheels of the Labor Market?” *Brookings Papers on Economic Activity*, 2001(1): 69–133.
- Bound, John, and Harry J. Holzer.** 2000. “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s.” *Journal of Labor Economics*, 18(1): 20–54.
- Cadena, Brian C., and Brian K. Kovak.** 2016. “Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession.” *American Economic Journal: Applied Economics*, 8(1): 257–290.
- Card, David.** 2001. “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration.” *Journal of Labor Economics*, 19(1): 22–64.
- Coen-Pirani, Daniele.** 2010. “Understanding Gross Worker Flows across US States.” *Journal of Monetary Economics*, 57(7): 769–784.
- Dao, Mai, Davide Furceri, and Prakash Loungani.** 2017. “Regional Labor Market

- Adjustment in the United States: Trend and Cycle.” *Review of Economics and Statistics*, 99(2): 243–257.
- Decressin, Jörg, and Antonio Fatás.** 1995. “Regional Labor Market Dynamics in Europe.” *European Economic Review*, 39(9): 1627–1655.
- Diamond, Rebecca.** 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000.” *American Economic Review*, 106(3): 479–524.
- Dix-Carneiro, Rafael, and Brian K. Kovak.** 2017. “Trade Liberalization and Regional Dynamics.” *American Economic Review*, 107(10): 2908–46.
- Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy.** 2010. “Housing Busts and Household Mobility.” *Journal of Urban Economics*, 68(1): 34–45.
- Foote, Andrew.** 2016. “The Effects of Negative House Price Changes on Migration: Evidence across US Housing Downturns.” *Regional Science and Urban Economics*, 60: 292–299.
- Foster, Lucia, John Haltiwanger, and Chad Syverson.** 2008. “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review*, 98(1): 394–425.
- Gathmann, Christina, Ines Helm, and Uta Schönberg.** 2016. “Spillover Effects in Local Labor Markets: Evidence from Mass Layoffs.” <http://www.utaschoenberg.com/working-papers>.
- Glaeser, Edward L.** 2005. “Urban Colossus: Why is New York America’s Largest City?” NBER Working Paper No. 11398.
- Glaeser, Edward L., and Giacomo A.M. Ponzetto.** 2007. “Did the Death of Distance Hurt Detroit and Help New York?” NBER Working Paper No. 13710.
- Glaeser, Edward L., and Jesse Shapiro.** 2001. “Is There a New Urbanism? The Growth of US cities in the 1990s.” NBER Working Paper No. 8357.
- Glaeser, Edward L., and Joseph Gyourko.** 2005. “Urban Decline and Durable Housing.” *Journal of Political Economy*, 113(2): 345–000.
- Glaeser, Edward L., and Joshua D. Gottlieb.** 2009. “The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States.” *Journal of Economic Literature*, 47(4): 983–1028.
- Glaeser, Edward L., Jed Kolko, and Albert Saiz.** 2001. “Consumer City.” *Journal of Economic Geography*, 1(1): 27–50.
- Glaeser, Edward L., José A. Scheinkman, and Andrei Shleifer.** 1995. “Economic

- Growth in a Cross-Section of Cities.” *Journal of Monetary Economics*, 36(1): 117–143.
- Glaeser, Edward L., Joseph Gyourko, and Raven E. Saks.** 2006. “Urban Growth and Housing Supply.” *Journal of Economic Geography*, 6(1): 71–89.
- Glaeser, Edward L., Sari Pekkala Kerr, and William R. Kerr.** 2015. “Entrepreneurship and Urban Growth: An Empirical Assessment with Historical Mines.” *Review of Economics and Statistics*, 97(2): 498–520.
- Goldin, Claudia D., and Lawrence F. Katz.** 2008. *The Race between Education and Technology*. Harvard University Press.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2017. “Bartik Instruments: What, When, Why, and How.” <https://sites.google.com/site/isaacsorkin/papers>.
- Gonzalez, Arturo.** 1998. “Mexican Enclaves and the Price of Culture.” *Journal of Urban Economics*, 43(2): 273–291.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti.** 2010. “Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings.” *Journal of Political Economy*, 118(3): 536–598.
- Hornbeck, Richard.** 2012. “The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe.” *American Economic Review*, 102(4): 1477–1507.
- Howard, Greg.** 2017. “The Migration Accelerator: Labor Mobility, Housing, and Aggregate Demand.” <https://sites.google.com/view/greghoward>.
- Jaeger, David A., Joakim Ruist, and Jan Stuhler.** 2017. “Shift-Share Instruments and the Impact of Immigration.” <https://janstuhler.wordpress.com/research>.
- Jimeno, Juan F., and Samuel Bentolila.** 1998. “Regional Unemployment Persistence (Spain, 1976–1994).” *Labour Economics*, 5(1): 25–51.
- Jordà, Òscar.** 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review*, 95(1): 161–182.
- Kaplan, Greg, and Sam Schulhofer-Wohl.** 2017. “Understanding the Long-Run Decline in Interstate Migration.” *International Economic Review*, 58(1): 57–94.
- Kennan, John, and James R. Walker.** 2011. “The Effect of Expected Income on Individual Migration Decisions.” *Econometrica*, 79(1): 211–251.
- Kline, Patrick, and Enrico Moretti.** 2013. “Place Based Policies with Unemployment.” *American Economic Review*, 103(3): 238–243.
- Kline, Patrick, and Enrico Moretti.** 2014. “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley

- Authority.” *Quarterly Journal of Economics*, 129(1): 275–331.
- Koo, Jahyeong, Keith R. Phillips, and Fiona D. Sigalla.** 2000. “Measuring Regional Cost of Living.” *Journal of Business and Economic Statistics*, 18(1): 127–136.
- Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Rugles.** 2017. “IPUMS National Historical Geographic Information System: Version 12.0 [Database].” Minneapolis: University of Minnesota.
- Michaels, Guy, Ferdinand Rauch, and Stephen J. Redding.** 2012. “Urbanization and Structural Transformation.” *Quarterly Journal of Economics*, 127(2): 535–586.
- Molloy, Raven, Christopher L. Smith, and Abigail Wozniak.** 2011. “Internal Migration in the United States.” *Journal of Economic Perspectives*, 25(3): 173–96.
- Molloy, Raven, Christopher L. Smith, and Abigail Wozniak.** 2017. “Job Changing and the Decline in Long-Distance Migration in the United States.” *Demography*, 54(2): 631–653.
- Monras, Joan.** 2015a. “Economic Shocks and Internal Migration.” IZA Discussion Paper No. 8840.
- Moretti, Enrico.** 2004. “Human Capital Externalities in Cities.” In *Handbook of Urban and Regional Economics*. Vol. 4, , ed. J. Vernon Henderson and Jacques F. Thisse, 2243–2291. Amsterdam: Elsevier.
- Moretti, Enrico.** 2010. “Local Multipliers.” *American Economic Review*, 100(2): 373–377.
- Moretti, Enrico.** 2011. “Local Labor Markets.” In *Handbook of Labor Economics*. Vol. 4B, , ed. David Card and Orley Ashenfelter, 1237–1313. New York: Elsevier.
- Moretti, Enrico.** 2012. *The New Geography of Jobs*. New York: Houghton Mifflin Harcourt.
- Munshi, Kaivan.** 2003. “Networks in the Modern Economy: Mexican Migrants in the US Labor Market.” *Quarterly Journal of Economics*, 118(2): 549–599.
- Nickell, Stephen.** 1981. “Biases in Dynamic Models with Fixed Effects.” *Econometrica*, 49(6): 1417–1426.
- Notowidigdo, Matthew J.** 2011. “The Incidence of Local Labor Demand Shocks.” NBER Working Paper No. 17167.
- Obstfeld, Maurice, and Giovanni Peri.** 1998. “Regional Non-Adjustment and Fiscal Policy.” *Economic Policy*, 13(26): 207–259.
- OECD.** 2005a. “How Persistent Are Regional Disparities in Employment? The Role of Geographic Mobility.” OECD Employment Outlook.
- Oi, Walter Y.** 1996. “The Welfare Implications of Invention.” In *The Economics of New Goods*. , ed. Timothy F. Bresnahan and Robert J. Gordon, 109–142. Chicago: University

of Chicago Press.

- Overman, Henry G., and Diego Puga.** 2002. “Unemployment Clusters across Europe’s Regions and Countries.” *Economic Policy*, 17(34): 115–148.
- Partridge, Mark D., Dan S. Rickman, M. Rose Olfert, and Kamar Ali.** 2012. “Dwindling U.S. Internal Migration: Evidence of Spatial Equilibrium or Structural Shifts in Local Labor Markets?” *Regional Science and Urban Economics*, 42(1): 375–388.
- Phillips, Keith R., and Christina Daly.** 2010. “Improving the ACCRA US Regional Cost of Living Index.” *Journal of Economic and Social Measurement*, 35(1): 33–42.
- Rappaport, Jordan.** 2007. “Moving to Nice Weather.” *Regional Science and Urban Economics*, 37(3): 375–398.
- Rappaport, Jordan.** 2012. “Why Does Unemployment Differ Persistently Across Metro Areas?” *Federal Reserve Bank of Kansas City, Economic Review*, 97(2): 5–35.
- Rappaport, Jordan, and Jeffrey D. Sachs.** 2003. “The United States as a Coastal Nation.” *Journal of Economic Growth*, 8(1): 5–46.
- Roback, Jennifer.** 1982. “Wages, Rents, and the Quality of Life.” *Journal of Political Economy*, 90(6): 1257–1278.
- Rosen, Sherwin.** 1979. “Current Issues in Urban Economics.” , ed. Peter N. Miezowski and Mahlon R. Straszheim, Chapter Wage-based Indexes of Urban Quality of Life, 74–104. Baltimore: Johns Hopkins University Press.
- Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek.** 2017. “Integrated Public Use Microdata Series: Version 7.0 [dataset].” Minneapolis: University of Minnesota. Minneapolis: University of Minnesota.
- Saiz, Albert.** 2010. “The Geographic Determinants of Housing Supply.” *Quarterly Journal of Economics*, 125(3): 1253–1296.
- Sinai, Todd, and Nicholas Souleles.** 2013. “Can Owning a Home Hedge the Risk of Moving?” *American Economic Journal: Economic Policy*, 5(2): 282–312.
- Sinai, Todd, and Nicholas S. Souleles.** 2005. “Owner-Occupied Housing as a Hedge against Rent Risk.” *Quarterly Journal of Economics*, 120(2): 763–789.
- Spring, Amy, Stewart E. Tolnay, and Kyle Crowder.** 2016. “Moving for Opportunities? Changing Patterns of Migration in North America.” In *International Handbook of Migration and Population Distribution*. Vol. 6, , ed. Michael J. White, 421–448. New York: Springer.
- Stock, James H., and Motohiro Yogo.** 2005. “Testing for Weak Instruments in Linear IV Regression.” In *Identification and Inference for Econometric Models: Essays in Honor*

- of *Thomas Rothenberg*, ed. Donald W. K. Andrews and James H. Stock, 80–108. New York: Cambridge University Press.
- Tinbergen, Jan.** 1974. “Substitution of Graduate by Other Labour.” *Kyklos*, 27(2): 217–226.
- Tolbert, Charles M., and Molly Sizer.** 1996. “U.S. Commuting Zones and Labor Market Areas: A 1990 Update.” Economic Research Service Staff Paper No. 9614.
- Wozniak, Abigail.** 2010. “Are College Graduates More Responsive to Distant Labor Market Opportunities?” *Journal of Human Resources*, 45(4): 944–970.
- Yagan, Danny.** 2017. “Employment Hysteresis from the Great Recession.” NBER Working Paper No. 23844.