

The Incidence of Local Labor Demand Shocks

Matthew J. Notowidigdo, *University of Chicago Booth School of Business and National Bureau of Economic Research*

Low-skill workers are comparatively immobile. This paper estimates the role of housing prices and social transfers in accounting for this fact using a spatial equilibrium model. Reduced-form estimates using US census data show that positive local labor demand shocks increase population more than negative shocks reduce population, that this asymmetry is larger for low-skill workers, and that such an asymmetry is absent for average wages, housing values, and rental prices. Generalized method of moments estimates reveal that the comparative immobility of low-skill workers is due not to higher mobility costs but to a lower incidence of adverse labor demand shocks.

I. Introduction

When a city experiences an adverse labor demand shock, the share of the adult population with a college degree tends to decline (Glaeser and Gyourko

I am grateful to Daron Acemoglu, David Autor, and Amy Finkelstein for their guidance and support. I thank Ting Wang for helpful research assistance in preparing the final version of the manuscript. I also thank Leila Agha, Josh Angrist, Fernando Duarte, Tal Gross, Erik Hurst, Cynthia Kinnan, Jean-Paul L'Huillier, Amanda Pallais, Jim Poterba, Michael Powell, Nirupama Rao, and Bill Wheaton as well as seminar participants at Massachusetts Institute of Technology (MIT), Northwestern, Berkeley, Harvard Kennedy School, University of Chicago, Wharton, Stanford, Yale School of Management, Brown, and the 2011 Society for Economic Dynamics (SED) Annual Meeting for helpful comments. I gratefully acknowledge the National Institute of Aging (grant T32-AG000186) and the MIT Shultz Fund for financial support. Contact the author at noto@chicagobooth.edu. Information concerning access to the data used in this paper is available as supplemental material online.

[*Journal of Labor Economics*, 2020, vol. 38, no. 3]

© 2020 by The University of Chicago. All rights reserved. 0734-306X/2020/3803-0002\$10.00
Submitted December 16, 2016; Accepted May 7, 2019; Electronically published May 27, 2020

2005). A standard explanation for this pattern is that barriers to mobility are greater for low-skill workers (Topel 1986, S183; Bound and Holzer 2000).¹

This paper proposes and tests an alternative explanation that focuses on why low-skill workers may be disproportionately compensated during adverse labor demand shocks rather than why it may be disproportionately costly for them to out-migrate. This explanation has two components. First, as documented below, adverse shocks substantially reduce the cost of housing. This fact and the existing evidence that the expenditure share on housing declines with income imply that low-skill workers are disproportionately compensated by housing price declines.² Second, means-tested public assistance programs disproportionately compensate low-skill workers during adverse shocks. I document below that, not surprisingly, aggregate transfer program expenditures are highly responsive to local labor market conditions.

These two different types of explanations—one based on mobility costs and one based on compensating factors—are not incompatible; however, their relative importance ultimately determines the actual incidence of local labor demand shocks. If out-migration of workers is low primarily because of mobility costs, then the incidence of local labor demand shocks will be primarily borne by workers; additionally, to the extent that mobility costs are greater for low-skill workers, they may disproportionately bear the incidence of the adverse shock. Alternatively, if the incidence of adverse local labor demand shocks is primarily borne by immobile housing and social insurance programs, then low-skill workers will be disproportionately compensated and, consequently, less likely to out-migrate.

In this paper, I develop and estimate a spatial equilibrium model that captures how wages, population, housing prices, and transfer payments reequilibrate following a shift in local labor demand. The model is based on the spatial equilibrium model in Roback (1982). Following Glaeser and Gyourko (2005), the model in this paper allows for a concave local housing supply

¹ The existence of greater barriers to mobility for low-skill workers is consistent with a large empirical literature that has documented that the local labor supply elasticity is larger for high-skill workers than for low-skill workers. For example, Bound and Holzer (2000) find that the elasticity of local labor supply with respect to wages is significantly higher for college-educated workers than for workers with no more than a high school education. Similarly, Topel (1986) finds that local labor demand shifts generate much smaller wage differentials among more educated workers. Topel writes that “consistent with the greater geographic mobility of more educated workers, their wages are less sensitive to both current and future changes in relative employment.”

² Of course, if low-skill workers are homeowners and not renters, then there is a negative wealth effect in addition to the decline in the user cost of housing following a negative local labor demand shock. Consistent with much of the recent urban economics literature (e.g., Glaeser and Gyourko 2005; Moretti 2013), I assume in the model below that everyone is a renter. I also explore alternative specifications that assume that the demand for housing is homothetic, so that the expenditure share on housing is assumed to be the same for high-skill and low-skill workers.

curve, arising from the durability of the local housing stock.³ While the Glaeser and Gyourko model assumes perfect mobility, I allow for heterogeneous mobility costs that limit spatial arbitrage, as in Topel (1986). Unlike the preceding models, I explicitly model local labor demand.

To give the basic intuition of the model, consider the following simplified version.⁴ The main conceptual experiment in the model is that a single city experiences a (positive or negative) labor demand shock while a large number of other cities remain unchanged. Figures 1 and 2 provide graphical representations of the different equilibrium responses of wages, population, and housing prices for four scenarios, depending on whether housing supply is constant elasticity or asymmetric and whether workers are perfectly mobile or face mobility costs when out-migrating.

Figure 1 depicts the equilibrium response when the elasticity of supply of housing is constant.⁵ The figure shows a positive shift in the labor demand curve, which raises wages by Δ . This increase in wages causes in-migration, which bids up housing prices until the increase in housing costs exactly offsets the wage increase (thus restoring the equilibrium no-arbitrage condition for workers). If workers are perfectly mobile, then the figure shows that the effect of a negative shock ($-\Delta$) is symmetric; that is, wages, housing prices, and population adjust by equal and opposite magnitudes (as shown by L^A in the figure). This symmetry comes from the log-linearity of the housing supply curve and the perfect mobility of workers. If, alternatively, workers face nonnegligible mobility costs, then there will be less out-migration following a negative shock. With nonnegligible mobility costs, the no-arbitrage condition is now that the marginal worker must be indifferent between staying and paying c to out-migrate. In this case, both the population and housing price responses are asymmetric: positive shocks increase population and housing prices more than negative shocks reduce them (see L^B in the figure). Intuitively, while mobility costs constrain out-migration, they do not similarly constrain in-migration because there are a large number of potential in-migrants with negligible mobility costs (since the single city is assumed to be small relative to the rest of the world). Therefore, the

³ Throughout the paper I use the term “concave housing supply curve” to imply that positive housing demand shocks increase housing prices less than equal-sized negative shocks reduce housing prices. More formally, a concave housing supply curve implies that $\partial^2(\text{housing price})/\partial(\text{housing supply})^2 < 0$.

⁴ In this simplified version of the model, workers in a city inelastically supply labor so that net migration fully determines local labor supply. Workers also do not differ in productivity, and there are no transfer payments. The full model below introduces high-skill and low-skill workers as well as transfer payments. Firms are perfectly mobile so that labor demand is perfectly elastic. Homogeneous housing units are supplied by absentee landlords who live in other cities, and workers consume a fixed expenditure share of housing (s_H).

⁵ This is equivalent to assuming that the housing supply curve is log linear.




FIG. 1.—Constant housing supply elasticity. This figure displays the equilibrium response when the housing supply elasticity is constant. The initial equilibrium wages, labor supply, and housing prices are given by the circle in the center of the figure. An exogenous increase in wages encourages in-migration until labor supply rises to L_+ . At this point, housing prices have risen to completely offset the increase in wages, restoring the no-arbitrage condition for workers. If there are no mobility costs, then the equilibrium response of an equal-sized exogenous decrease in wages is symmetric, as shown by L^A . If out-migration is costly, however, then following a negative shock, the marginal out-migrant must be indifferent between staying and paying c to out-migrate. These mobility costs cause both population and housing prices to respond asymmetrically: positive shocks increase population and housing prices more than negative shocks reduce them.

increase in population following a positive shock is the same whether or not workers face heterogeneous costs of out-migration (see L^{\pm} in the figure).

In figure 2, the housing supply elasticity is no longer constant. Specifically, housing is more elastically supplied following an increase in housing demand than a decrease in demand. As discussed in greater detail in the main text below and in section A.2 of the appendix (available online), this asymmetric housing supply curve is consistent with a simple model of durable housing where housing units are not destroyed once created (Glaeser and Gyourko 2005). When workers are perfectly mobile, housing prices respond symmetrically (despite the asymmetry in the housing supply curve). Intuitively, housing costs still must adjust to exactly offset the wage changes. Only population responds asymmetrically (as shown by L^c in the figure). However, if




FIG. 2.—Concave housing supply curve. This figure displays the equilibrium response when the housing supply curve is concave. As the main text and the appendix (available online) describe in more detail, a concave housing supply curve is consistent with a durable housing stock that is not destroyed once created. As in figure 1, the initial equilibrium wages, labor supply, and housing prices are given by the circle in the center of the figure. An exogenous increase in wages encourages in-migration until labor supply rises to L_+ . At this point, housing prices have risen to completely offset the increase in wages, restoring the no-arbitrage condition for workers. If there are no mobility costs, then housing prices still respond symmetrically (p^C). Intuitively, housing costs still must adjust to exactly offset the wage changes. Only population responds asymmetrically (as shown by L^C). If workers have mobility costs, then the asymmetry of the population response is even greater (see L^D), and in this case housing prices also respond asymmetrically.

workers have heterogeneous mobility costs to out-migrate as described above, then in this case the asymmetry of the population response is even greater (see L^D in the figure), and housing prices also respond asymmetrically.

These scenarios give the intuition for the following two implications of the full model derived below: (1) if positive labor demand shocks increase population more than negative shocks reduce population, this suggests the existence of a concave housing supply curve and/or heterogeneous mobility costs, and (2) if positive shocks increase housing prices more than negative shocks reduce housing prices, that is consistent with the existence of heterogeneous mobility costs. The full model below shows that these implications continue to hold in a richer setting with transfer payments and two types of workers.

The model guides the empirical strategy, which consists of two steps. In the first step, I test for asymmetric responses of wages, employment, population, and housing prices to symmetric labor demand shocks. The validity of this exercise requires constructing plausibly exogenous positive and negative shifts in local labor demand of equal magnitude. This paper follows Bartik (1991) in constructing an instrumental variable for local labor demand shocks by interacting cross-sectional differences in industrial composition with national changes in industry employment shares. I find robust evidence using US census data that positive local labor demand shocks increase population (and employment) more than negative shocks reduce population (and employment) and that this asymmetry is greater for low-skill workers. These robust asymmetric relationships for local population and employment contrast sharply with the absence of any evidence of a similar asymmetric relationship for (any measure of) wages, housing values, and rental prices, although all of these other variables respond strongly to local labor demand.⁶ As the spatial equilibrium model makes clear, these results are consistent with a concave local housing supply curve and limited mobility costs. While the Bartik (1991) procedure has been widely used in labor economics and urban economics, to my knowledge this is the first paper that uses this procedure to explicitly test for asymmetric responses of wages, employment, and population to local labor demand shocks.

To quantitatively estimate the magnitude of mobility costs by skill and the shape of the housing supply curve, in the second step of the empirical analysis I estimate the full spatial equilibrium model using a nonlinear simultaneous-equations generalized method of moments (GMM) estimator. The GMM estimates suggest that the housing supply curve is concave and that mobility costs (defined as a fraction of income) are at most modest and are comparable for both high-skill and low-skill workers. The GMM results reveal several other important findings. First, the observed asymmetric population responses are primarily accounted for by an asymmetric housing supply curve rather than by substantial barriers to mobility. Second, the

⁶ The model in Glaeser and Gyourko (2005) predicts a concave relationship between housing prices and the exogenous labor demand, and these authors find supportive evidence of this prediction using an exogenous shock based on climate. As discussed in more detail in sec. A.4 of the appendix, the key difference between the model in this paper and the model in Glaeser and Gyourko (2005) is that the model in this paper assumes that housing units are homogeneous, while in the Glaeser and Gyourko model housing units have heterogeneous, location-specific amenities. In other words, in the Glaeser and Gyourko model exogenous shocks induce compositional changes in the distribution of location-specific amenities in the housing stock, and these compositional changes affect the (unconditional) average housing price. The difference in empirical results comes from the fact that Glaeser and Gyourko (2005) use mean temperature to construct local amenity shocks based on a dummy variable for whether or not the January mean temperature is greater than 29.1°F, whereas I use variation in local labor demand.

results suggest that the observed difference in out-migration by skill is primarily accounted for by transfer payments rather than differences by skill in housing expenditure shares. Third, the results suggest that the primary explanation for the comparative immobility of low-skill workers is not higher mobility costs per se but rather a lower incidence of adverse local labor demand shocks. Consequently, much of the incidence of adverse labor demand shocks is diffused to homeowners, landlords, and public assistance programs.

Finally, I use the GMM estimates to construct counterfactual estimates of how local labor markets would adjust to shocks if the system of means-tested transfer payments was replaced with a system of mobility subsidies for both high-skill and low-skill workers. In this alternative system, the skill composition of the local labor force is much less responsive to shifts in local labor demand, but population continues to respond strongly asymmetrically because of the asymmetric housing supply curve. The estimation of the full model necessarily requires stronger assumptions than were needed to test for asymmetric responses to shocks. To be able to consistently estimate the relative magnitude of mobility costs by skill, I must assume that unobserved changes in local amenities induced by local labor demand shocks are not differentially valued by high-skill and low-skill workers. To be able to consistently estimate the absolute magnitude of mobility costs, however, a stronger assumption is needed—namely, that unobserved changes in local amenities are uncorrelated with local labor demand shocks. Because of this, the analysis of the absolute magnitudes of mobility costs should be interpreted more cautiously.

This paper is broadly related to recent empirical work that acknowledges the importance of migration costs in determining spatial equilibrium. This work has emphasized the importance of imperfect mobility in determining the efficiency of place-based policies (Busso, Gregory, and Kline 2013) and in determining the marginal willingness to pay for environmental amenities (Bayer, Keohane, and Timmins 2008). This paper is also related to recent work on the effects of wage income and welfare income on the individual migration decision (Kennan and Walker 2010, 2011); this paper is highly complementary to these two papers, which employ a very different empirical approach by estimating a rich structural model of individual migration decisions.⁷

The rest of the paper proceeds as follows. Section II presents the theoretical framework. Section III discusses the empirical strategy and the data. Section IV presents the reduced-form empirical results. Section V presents GMM estimates of the full model. Section VI concludes.

II. Theoretical Framework

This section presents a simple spatial equilibrium model of a local labor market that captures how wages, population, housing prices, and transfer

⁷ Also related to this paper is the recent literature on the causal effect of education and geographic mobility (Wozniak 2010; Malamud and Wozniak 2012).

payments reequilibrate following a local labor demand shock.⁸ The heart of the model is a no-arbitrage condition in which the marginal worker is indifferent between remaining in the city receiving the shock and moving away (Roback 1982). This condition implicitly defines a local labor supply curve that determines the amount of migration in response to a labor demand shock. The model below allows for mobility costs, which limit spatial arbitrage and cause the incidence of the labor demand shock to fall at least partially on workers (Topel 1986).⁹ Additionally, the model admits two types of workers (high skill and low skill) who differ in productivity, imperfectly substitute in production, and may potentially differ in their housing expenditure shares, eligibility for transfer payments, and mobility costs. If an adverse labor demand shock causes relatively greater out-migration of high-skill labor, the model clarifies when this is because the incidence of the shock is borne by other factors that disproportionately compensate low-skill workers and when this is due to greater barriers to mobility for low-skill workers.

The conceptual experiment is that a single city (out of a large universe of cities) experiences a labor demand shock between the first and second period. For simplicity, the model is presented as a 2-period model in order to rule out the effects of long-run expectations, the differences between temporary and permanent shocks, the option value from moving, and other issues arising in dynamic spatial equilibrium models. I focus on decadal changes in the empirical analyses below in order to minimize the influence of these other factors, and I leave a rigorous treatment of these dynamics for future work.

To give the general intuition of the model, consider an adverse local labor demand shock in a city. This shock will reduce wages, which encourages out-migration and, ultimately, lowers housing prices until the no-arbitrage condition is restored for the marginal worker. The amount of out-migration is determined by the magnitude of mobility costs, the generosity of transfer payments, and the elasticity of supply of housing in response to a decline in housing demand.

The four main components of the model (labor demand, transfer payments, housing market, and labor supply) are now discussed in detail.

A. Labor Demand

Assume a large number of cities indexed by i , and define the (large) number of high-skill and low-skill workers in city i and time t as H_{it} and L_{it} . Production

⁸ The model is a “local general equilibrium” model in the sense that labor demand shocks affect nonlabor markets within the city; however, it is not a full general equilibrium model because when the single city is shocked, the (minimal) effects on the rest of the universe are ignored.

⁹ Topel (1986) is primarily concerned with understanding differences between permanent and transitory shocks; in the simple 2-period model in this paper, all shocks are necessarily permanent.

of the homogeneous tradable good y is given by the following constant elasticity of substitution aggregate production function:¹⁰

$$y_{it} = \theta_{it}((1 - \lambda)L_{it}^\rho + \lambda(\zeta H_{it})^\rho)^{\alpha/\rho},$$

where λ is a share parameter, α measures the returns to scale of the labor aggregate, ζ is the relative efficiency of high-skill labor, and ρ is related to the elasticity of substitution between high-skill and low-skill labor by $\sigma_{H,L} \equiv 1/(1 - \rho)$.¹¹ The θ_{it} term is a city-specific index of local labor demand. In the empirical section below, I argue that my instrumental variable for local labor demand is a plausibly exogenous source of variation in θ_{it} .

Assuming that wages are set on the demand curve, then they are given by the following marginal productivity conditions:

$$\begin{aligned} w_{it}^H &= \alpha \theta_{it} ((1 - \lambda)L_{it}^\rho + \lambda(\zeta H_{it})^\rho)^{(\alpha-\rho)/\rho} \lambda \zeta (\zeta H_{it})^{\rho-1}, \\ w_{it}^L &= \alpha \theta_{it} ((1 - \lambda)L_{it}^\rho + \lambda(\zeta H_{it})^\rho)^{(\alpha-\rho)/\rho} (1 - \lambda) (L_{it})^{\rho-1}. \end{aligned}$$

Totally differentiating the above wage expressions results in the following conditions for the evolution of wages in terms of exogenous labor demand shock ($\Delta\theta_{it}$) and the endogenous migration responses (ΔH_{it} and ΔL_{it}):

$$\Delta w_{it}^H = \Delta\theta_{it} + ((\rho - 1) + (\alpha - \rho)(\pi))\Delta H_{it} + (\alpha - \rho)(1 - \pi)\Delta L_{it}, \quad (1)$$

$$\Delta w_{it}^L = \Delta\theta_{it} + ((\rho - 1) + (\alpha - \rho)(1 - \pi))\Delta L_{it} + (\alpha - \rho)(\pi)\Delta H_{it}, \quad (2)$$

where $\pi = \lambda(\zeta H)^\rho / ((1 - \lambda)L^\rho + \lambda(\zeta H)^\rho)$ and the Δ operator represents the percentage change over time.

B. Transfer Payments

Means-tested public assistance programs are available only to low-skill workers and are modeled as a constant elasticity function of wages, $b_{it} = \bar{B} \cdot (w_{it}^L)^\Psi$, where b_{it} is the transfer income (social assistance benefits) for the representative low-skill worker, \bar{B} is a constant, and Ψ is the elasticity of public assistance income with respect to low-skill wages.¹² The constant

¹⁰ For simplicity, capital is not included in the model. This could be important if part of the incidence of labor demand shocks falls on owners of capital. Since the empirical results are based on decadal changes, it seems reasonable to assume that the elasticity of supply of capital over this time period is fairly large.

¹¹ Let μ be the share of high-skill workers in the labor market. Then if $\lambda = (1 - \mu)^{\rho-1} / ((\zeta\mu)^{\rho-1} + (1 - \mu)^{\rho-1})$, ζ will give the equilibrium wage premium.

¹² Using Panel Study of Income Dynamics (PSID) data from 1990, I calculate that 0.5% of households receiving Aid to Families with Dependent Children (AFDC) income during the past year had a household head with at least a college degree. Among households receiving food stamps during the past year, the fraction is 0.7%. By contrast, among households receiving AFDC income, 79.1% had a household head with a high school education or less; for food stamps, the fraction

elasticity assumption is a simplification; empirically, I do not find robust evidence of a nonlinear or asymmetric effect of labor demand shocks on aggregate expenditures on transfer programs, so this assumption appears to be a reasonable approximation. The equations above imply the following expression for the evolution of transfer income in response to changes in low-skill wages:

$$\Delta b_{it} = \Psi \Delta w_{it}^L. \quad (3)$$

I assume $\Psi < 0$, which implies that transfer programs provide wage insurance, and I define s_b^L as the share of total income that comes from transfer program benefits for low-skill workers; for high-skill workers, $s_b^H = 0$ by assumption.

C. Housing Market

Workers consume housing and a tradable consumption good. A homogeneous housing stock is supplied by absentee landlords, and the aggregate housing supply curve is given by $\mathcal{H}^S(p_{it}^H)$, where p_{it}^H is the price of housing. Defining $s_{\mathcal{H}}^H$ and $s_{\mathcal{H}}^L$ as the housing expenditure shares for high-skill and low-skill workers, respectively, the main GMM estimates assume that $s_{\mathcal{H}}^L > s_{\mathcal{H}}^H$. This assumption is consistent with observed differences in housing expenditure shares by level of education (which here is a proxy for permanent income). I also report results that assume $s_{\mathcal{H}}^L = s_{\mathcal{H}}^H$ in robustness analysis below.

Instead of assuming a specific utility function to derive the demand for housing, I instead approximate aggregate demand for housing as follows:

$$\mathcal{H}^D(p_{it}^H) = \frac{s_{\mathcal{H}}^H \omega_{it}^H H_{it} + s_{\mathcal{H}}^L (b_{it}^L + \omega_{it}^L) L_{it}}{p_{it}^H}.$$

This expression is an approximation, since I am implicitly assuming that any changes in income induced by a shift in labor demand are small so that income effects can be ignored. Empirically, the changes in wages within skill groups are small relative to the differences in wages across skill groups, so this assumption is sensible.

The initial supply-demand equilibrium in housing market in the first period is given by $\mathcal{H}^S(p_{it}^H) = \mathcal{H}^D(p_{it}^H)$. Totally differentiating this equilibrium condition gives the following expression for the housing market response:

$$\Delta p_{it}^H + \Delta \mathcal{H}^S(\Delta p_{it}^H) = \nu(\Delta y_{it}^H + \Delta H_{it}) + (1 - \nu)(\Delta y_{it}^L + \Delta L_{it}), \quad (4)$$

is 82.6%. Therefore, means-tested public assistance program benefits are primarily used by households with low-skill household heads.

where ν is the high-skill share of aggregate housing demand and Δy_{it}^j gives the change in total income for skill group $j(\in \{H, L\})$; that is, $\Delta y_{it}^j = s_b^j \Delta b_{it}^j + (1 - s_b^j) \Delta w_{it}^j$. If the housing supply curve has constant elasticity, then $\Delta H^S(p_{it}^H) = \sigma \cdot \Delta p_{it}^H$. Since housing is a durable good, however, the housing supply elasticity is not likely to be constant. Instead, the housing supply elasticity will be larger for increases in housing demand than for decreases in housing demand because of the durability of the housing stock (Glaeser and Gyourko 2005). Formally, durable housing implies that $\Delta \mathcal{H}^S(\Delta p_{it}^H)$ is convex in Δp_{it}^H . Section A.4 of the appendix presents a simple model that provides microfoundations for a concave housing supply curve based on slow depreciation of the housing stock and a heterogeneous distribution of costs of supplying housing.

D. Labor Supply

For simplicity, I assume that workers inelastically supply labor to their local labor market, so that all variation in local employment comes only from migration decisions. The local labor supply curve is then implicitly defined by a mobility condition that states that the marginal migrant must be indifferent between remaining in city i and moving to any other city.

I introduce costly spatial arbitrage by assuming that workers have heterogeneous mobility costs. I construe mobility costs broadly to encompass both financial and psychic barriers to out-migration as well as heterogeneous tastes and distastes for a given location. Thus, unlike Topel (1986) I allow mobility costs to take on positive and negative values. Positive values encompass both actual moving costs and preferences for the current city, while negative values represent distaste of potential in-migrants for a given area. Formally, I model this by assuming that mobility costs for workers in city i are independently drawn from distributions $M_i^H(m)$ and $M_i^L(m)$ (with support $[0, \infty)$), while the mobility costs of in-migrating into city i for the workers living in all of the other cities are drawn from the distributions $M_{-i}^H(m)$ and $M_{-i}^L(m)$ (with support $(0, \infty]$). Mobility costs are defined as a fraction of total income, so that the marginal migrant receiving $(w + b)$ in city i will pay $(w + b)m$ to out-migrate. These mobility cost distributions imply mobility cost functions $c^H(\Delta H_{it})$ and $c^L(\Delta L_{it})$, which return the mobility cost of the marginal migrant given the change in population between the first and second period. For a smooth distribution of mobility costs, the mobility cost function will be strictly decreasing, so that the mobility cost of the marginal migrant increases as more workers out-migrate.¹³

¹³ Note that this 2-period model contains two important simplifications that make it straightforward to study mobility costs. First, following Topel (1986), gross migration will always equal net migration, so that there is only one marginal migrant per worker type in each city. The work of Artuç, Chaudhari, and McLauren (2010) and Chaudhari and McLauren (2007) suggest a tractable way

To derive the (implicit) labor supply curve for low-skill workers, let $v_i(\omega_{ii}^L + b_{ii}^L, p_{ii}^H)$ be the indirect utility function for the marginal low-skill worker in city i . Spatial equilibrium in the first period requires that the following condition holds for the marginal low-skill migrant in city i :

$$v_i(\omega_{ii}^L + b_{ii}^L, p_{ii}^H) = v_j(\omega_{ji}^L + b_{ji}^L, p_{ji}^H) \quad \forall j \neq i.$$

Now consider a shock to θ_i in city i . The shock will cause a wage differential that will encourage costly migration to arbitrage the wage and employment differential, and the price of housing and transfer payments will also adjust as a local general equilibrium response to the shock. Differentiating the above spatial equilibrium condition and applying Roy's identity results in the following expression:¹⁴

$$(1 - s_b^L)\Delta\omega_{ii}^L + s_b^L\Delta b_{ii} - s_{p_{ii}^H}^L\Delta p_{ii}^H + c^L(\Delta L_{ii}) = 0, \tag{5}$$

where $s_b^L (= b^L/(\omega^L + b^L))$ is public assistance income as a share of total income. An analogous expression holds for high-income workers (where $s_b^H = 0$):

$$\Delta\omega_{ii}^H - s_{p_{ii}^H}^H\Delta p_{ii}^H + c^H(\Delta H_{ii}) = 0. \tag{6}$$

Equations (5) and (6) are implicit labor supply curves because net migration is determined by the spatial equilibrium condition for the marginal migrant. In other words, the conditions above state that the change in indirect utility in response to changes in wages, transfer payments, and housing prices must equal the mobility costs of the marginal migrants. The ΔL_{ii} and ΔH_{ii} terms represent the amount of net migration that needs to occur to make these two equations hold.

to relax this assumption and allow gross migration flows to exceed net migration flows. Second, the mobility cost function is allowed to be asymmetric, but since this is a 2-period model the shape of this function does not depend on the history of past shocks. In a fully dynamic model, the history of past shocks may affect the elasticity of supply of in-migrants and out-migrants.

¹⁴ The full derivation of eq. (5) is given below: following a shock to city i , the new spatial equilibrium following the shock will be given by the following expression:

$$\frac{dv_i}{d \log(\theta_i)} + \frac{\partial v}{\partial(\omega + b)} m(\Delta L_{ii}) = 0.$$

In other words, this means that the change in indirect utility to the marginal migrant must equal that migrant's mobility costs, scaled by the marginal utility of (total) income. The argument ΔL_{ii} is the equilibrium change in population in response to the shock; i.e., $d \log(L_{ii})/d \log \theta_i$, which is used to "pick out" the marginal migrant after the population has changed by ΔL_{ii} . Computing the full derivative $dv_i/d\theta_i$ and applying Roy's identity yields eq. (5).

These two equations highlight the three reasons discussed in the introduction why net migration rates may differ by skill. First, public assistance program benefits are means tested, so that $s_b^L > s_b^H$. Second, low-skill workers consume a larger fraction of their income on housing $s_H^L > s_H^H$, meaning that housing price declines disproportionately compensate low-skill workers. Finally, the mobility cost functions may differ by skill. If low-skill workers typically face higher mobility costs following a negative shock, then $c^L(x) > c^H(x) \quad \forall x < 0$.

E. Equilibrium

Following an exogenous shock to local labor demand ($\Delta\theta_{it}$), the new equilibrium of the model is defined by the following conditions:

- Labor demand adjusts so that high-skill and low-skill wages equal marginal products (eqq. [1], [2]).
- Transfer payments adjust according to changes in low-skill wages (eq. [3]).
- Housing prices adjust so that the change in housing demand equals the change in housing supply (eq. [4]).
- Population adjusts so that the marginal high-skill and low-skill migrant is indifferent between staying and leaving (eqq. [5], [6]).

Although the nonlinearities in the housing supply curve ($\Delta H^S(\Delta p_{it}^H)$) and the mobility cost functions ($c^H(\Delta H_{it})$ and $c^L(\Delta L_{it})$) preclude analytical solutions without particular functional form assumptions, section A.2 of the appendix derives comparative statics for specific scenarios under the special case of constant returns to scale of production ($\alpha = 1$). Figure 3 reports results from simulating the model, with the details of the simulation given in section A.3 of the appendix. The figure shows that if population responds asymmetrically, it suggests the existence of a concave housing supply curve and/or the existence of heterogeneous mobility costs. The responsiveness of housing prices isolates the importance of heterogeneous mobility costs, since mobility costs cause immobile workers to bid up the price of housing during negative shocks, causing housing prices to respond asymmetrically. Therefore, the model suggests that it is possible to identify both mobility costs and the shape of the housing supply curve by using information on the joint responses of wages, population, housing prices, and transfer payments to exogenous labor demand shocks.

These simulations motivate the two-part empirical strategy below. First, I will estimate nonlinear reduced-form regressions to test for asymmetric responses to labor demand shocks. Second, I will carry out a full estimation of the model to recover the parameters that govern the distribution of mobility costs and the shape of the housing supply curve.

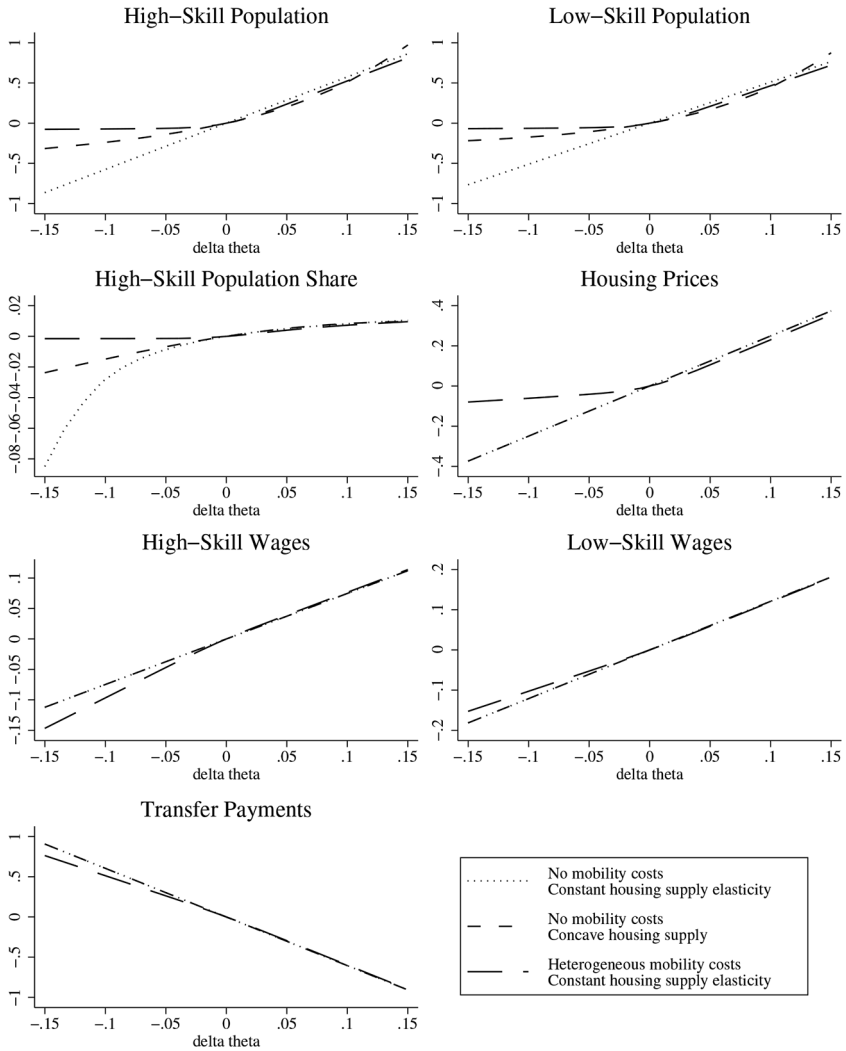


FIG. 3.—Model simulations. This figure displays simulated data from the model described in section II. See the appendix (available online) for more details on the simulation. The graphs clarify that an asymmetric response of population to the labor demand shock (“ $\Delta\theta$ ”) indicates the existence of a concave housing supply curve and/or the existence of heterogeneous mobility costs. The response of housing prices isolates the importance of mobility costs.

III. Empirical Strategy and Data

As the model makes clear, the reduced-form relationships between each of the endogenous variables (Δw^H , Δw^L , ΔH , ΔL , Δp^H , Δp^L) and the labor demand shock $\Delta\theta$ are informative about the shape of housing supply curve and the presence of heterogeneous mobility costs. This motivates the following reduced-form estimating equation:

$$\Delta x_{it} = g^x(\Delta\theta_{it}) + \alpha_t + \Delta\varepsilon_{it},$$

where i indexes cities, t indexes time periods, x is one of the endogenous variables above, α_t captures proportional shocks to all cities in a given time period, ε_{it} is an error term, and $g(\cdot)$ is a function to be estimated. Nonparametric estimates of $g(\cdot)$ are reported graphically below. In addition to the nonparametric estimates, I also parameterize $g^x(\Delta\theta)$ as $\beta(\Delta\theta) + \delta(\Delta\theta)^2$, which leads to the following baseline reduced-form empirical specification that is reported in the tables:

$$\Delta x_{it} = \beta \cdot \Delta\theta_{it} + \delta \cdot (\Delta\theta_{it})^2 + \alpha_t + \Delta\varepsilon_{it}, \quad (7)$$

where x is the endogenous variable of interest, β and δ are the coefficients on a quadratic in $\Delta\theta_{it}$, and α_t are year fixed effects. This reduced-form specification is estimated by ordinary least squares using a proxy for local labor demand (described below). The quadratic specification allows the elasticity of x_{it} with respect to θ_{it} to vary: specifically, the elasticity at $\Delta\theta_{it} = 0$ is given by $\hat{\beta}$, while $\hat{\beta} + 2\hat{\delta}\Delta\theta_{it}$ is the elasticity at $\Delta\theta_{it}$. Since the equation is estimated in first differences it implicitly controls for time-invariant differences across geographic areas, while the inclusion of year fixed effects captures any (proportional) changes in x_{it} common to all cities. Formally, the statistical test of $\delta \neq 0$ is sufficient to establish that positive and negative shifts in labor demand of equal magnitude have unequal effects. However, this test is evaluating the null hypothesis of a linear relationship against a specific parametric alternative. Therefore, I will also report nonparametric specification tests that test the null hypothesis of a linear relationship against a nonparametric alternative (Ellison and Ellison 2000).

Last, I also estimate the full model developed above to recover flexible estimates of the mobility cost functions of high-skill and low-skill workers and the housing supply curve parameters. The estimation is a nonlinear simultaneous-equations problem, and it is implemented using a two-step GMM estimator. The details of the GMM procedure are described in more detail below.

A. An Omnibus Instrumental Variable for Local Labor Demand

To estimate equation (7), a valid instrumental variable for local labor demand is needed. I follow the empirical strategy of Bartik (1991) and construct a measure of plausibly exogenous labor demand shocks derived by interacting cross-sectional differences in industrial composition with national changes

in industry employment shares.¹⁵ This relative demand index can be used to predict changes in wages and employment. The identifying assumption is that changes in industry shares at the national level are uncorrelated with city-level labor supply shocks and therefore represent plausibly exogenous (demand-induced) variation in metropolitan-area employment. This predicted employment variable (\hat{E}_{it}) is used to create a predicted change in local area employment ($\Delta\hat{\theta}_{it}$) as follows: $\Delta\hat{\theta}_{it} = (\hat{E}_{it} - E_{i,t-\tau})/E_{i,t-\tau}$. This measure is used as a proxy for $\Delta\theta_{it}$.¹⁶

The key identifying assumption is that this proxy is uncorrelated with unobserved shocks to local labor supply. In this paper a stronger assumption is also needed—specifically, I must assume that $\Delta\theta_{it} = X$ and $\Delta\theta_{it} = -X$ represent shifts in local labor demand of plausibly equal magnitude. This requirement gives a clear advantage to the Bartik procedure over other identifiable shocks to local labor demand, as this instrumental variable is an omnibus measure of changes in local labor demand. By contrast, if one were to use identifiable shifts to labor demand, such as movements in oil prices, coal prices, or other natural resource shocks, it would require that equal-sized positive and negative price changes represent equal-sized shifts in local labor demand. This may be difficult to justify in natural resource industries that are typically characterized by high amounts of specific capital and/or irreversible investments. An additional benefit of this procedure is that subsets of industries can be excluded when constructing the instrumental variable to verify that the results are not driven by particular sectors, which we investigate in the robustness analysis below.

An important piece of evidence in support of the key identifying assumption is that the distribution of the estimated labor demand shocks is highly symmetric (fig. A1, available online). This suggests that any estimated asymmetric responses is not being driven (in part) by an underlying asymmetric distribution of shocks.

B. Data and Descriptive Statistics

The data sources are briefly described here. Appendix section A.1 (“Data Appendix”) gives more detail on how the data set was created.

¹⁵ See Blanchard and Katz (1992), Bound and Holzer (2000), Luttmer (2005), and Autor and Duggan (2006) for other applications of this instrumental variable.

¹⁶ Formally, predicted employment growth is computed as follows:

$$\begin{aligned}\pi_{it} &= \sum_{k=1}^K \varphi_{i,k,t-\tau} \left(\frac{v_{-i,k,t} - v_{-i,k,t-\tau}}{v_{-i,k,t-\tau}} \right), \\ \hat{E}_{it} &= (1 + \pi_{it})E_{i,t-\tau}, \\ \Delta\hat{\theta}_{it} &= (\hat{E}_{it} - E_{i,t-\tau})/E_{i,t-\tau},\end{aligned}$$

where $\varphi_{i,k,t-\tau}$ is the employment share of industry k in city i and $v_{-i,k,t}$ is the national employment share of industry k excluding city i .

1. *Census Integrated Public Use Microdata Series (IPUMS)*

The basic panel of metropolitan-area data comes from the 1980, 1990, and 2000 census individual-level and household-level extracts from the IPUMS database (Ruggles et al. 2004).¹⁷ The baseline data are limited to individuals and households living in metropolitan areas. The IPUMS data are used to construct estimates of local area wages, employment, population, housing prices, and rental prices in each metropolitan area. The primary advantage of the census data is the ability to construct city-level measures disaggregated by skill. These data are also used to construct the predicted labor demand instrumental variable by using the industry categories of the individuals in the labor force. For remaining details, see appendix section A.1 (“Data Appendix”).

2. *Regional Economic Information System (REIS)*

The metropolitan-area measures of expenditures on public assistance programs are computed by aggregating the county-level aggregate data in the REIS. The REIS contains annual county-level data on total expenditures broken down by transfer program (e.g., food stamps, income maintenance programs, public medical benefits, veterans benefits, Supplemental Security Income [SSI] benefits). Counties are aggregated into metropolitan areas using the 1990 metropolitan statistical area (MSA) definitions. Because of the difficulty in aggregating counties into MSAs within Alaska and Virginia during this time period, MSAs in these states are dropped from the baseline sample. Although the data are not disaggregated below the county level, the data are based on government agency reports and are therefore quite reliable. According to recent work by Meyer, Mok, and Sullivan (2009), aggregate expenditure data may be sometimes preferable to individual or household survey data because of substantial underreporting in the latter.¹⁸ All transfer program measures are adjusted per low-skill capita based on the non-college-educated adult population.

Table 1 reports descriptive statistics for the final data set.

IV. Results

A. Graphical Evidence

Figures 4 and 5 report nonparametric reduced-form estimates for the primary dependent variables. In addition to the nonparametric estimates, linear

¹⁷ The 2007 American Community Survey is included as a robustness check. The 1970 census is not used at all because it identifies only a small subset of the MSAs that appear in later years.

¹⁸ Meyer, Mok, and Sullivan (2009) find substantial underreporting of benefit receipt for a wide range of transfer programs in a wide range of data sets, including the Current Population Survey, PSID, the Survey of Income and Program Participation, and the Consumer Expenditure Survey. They also document that the underreporting is not consistent over time.

Table 1
Summary Statistics

	N	Mean	SD	Percentile				
				5th	25th	50th	75th	95th
US census data (IPUMS):								
Adult population								
(in millions)	645	.425	.856	.060	.093	.177	.392	1.477
Employment (in millions)	645	.303	.596	.041	.067	.127	.283	1.036
Employment-to-population ratio	645	.711	.051	.625	.680	.714	.748	.786
Income per adult (in \$000s)	645	14.979	3.167	10.516	12.871	14.664	16.674	20.079
Residualized wage (\$)	645	11.528	1.198	9.790	10.710	11.387	12.275	13.685
Residualized wage, LFP adjusted (\$)	645	8.212	1.124	6.605	7.485	8.136	8.896	10.042
College share of adult population	645	.190	.063	.105	.143	.181	.226	.305
College share of employment	645	.221	.065	.131	.173	.213	.257	.341
Average housing value (in \$000s)	645	97.449	45.450	58.005	71.527	84.774	107.212	196.809
Average gross rent (in \$000s)	645	5.229	1.014	4.055	4.579	5.017	5.581	7.196
REIS data:								
Food stamps + income maintenance (in \$000s per non-college-educated adult)	645	.652	.325	.247	.429	.594	.792	1.286

NOTE.—The baseline sample is a balanced panel of 215 metropolitan statistical areas (MSAs), and all observations are MSA-year. Integrated Public Use Microdata Series (IPUMS) data are from the 1980, 1990, and 2000 census extracts. The Regional Economic Information System (REIS) data are county level and annual but are aggregated to MSAs using the 1990 MSA definitions. All dollar values in this table are nominal, but all dollar-valued variables are converted to real dollars in the analysis. All specifications in subsequent tables are in first differences, so the three decades in this data set become two 10-year changes. LFP = labor force participation.

estimates are graphed for comparison. The figures also display bootstrapped (uniform) 95% confidence intervals.¹⁹ The confidence intervals are very wide at the extremes, making it difficult to reject the null hypothesis that the data are described by any linear relationship. However, in some cases the confidence intervals reject the specific linear relationship estimated using a parametric linear model, although this visual test ignores estimation error in the linear model. Consequently, the nonparametric specification tests reported

¹⁹ The bootstrapped confidence intervals are computed on the basis of 10,000 replications, where MSAs are sampled with replacement. In each bootstrap step, an undersmoothed local linear bandwidth is chosen following Hall (1992). That paper reports Monte Carlo results that suggest that undersmoothing produces confidence interval estimates with greater coverage accuracy than confidence intervals obtained by explicit bias correction. The bandwidth of the Epanechnikov kernel used for point estimation is 0.041; the undersmoothed kernel bandwidth is $0.75 \times 0.041 = 0.031$.

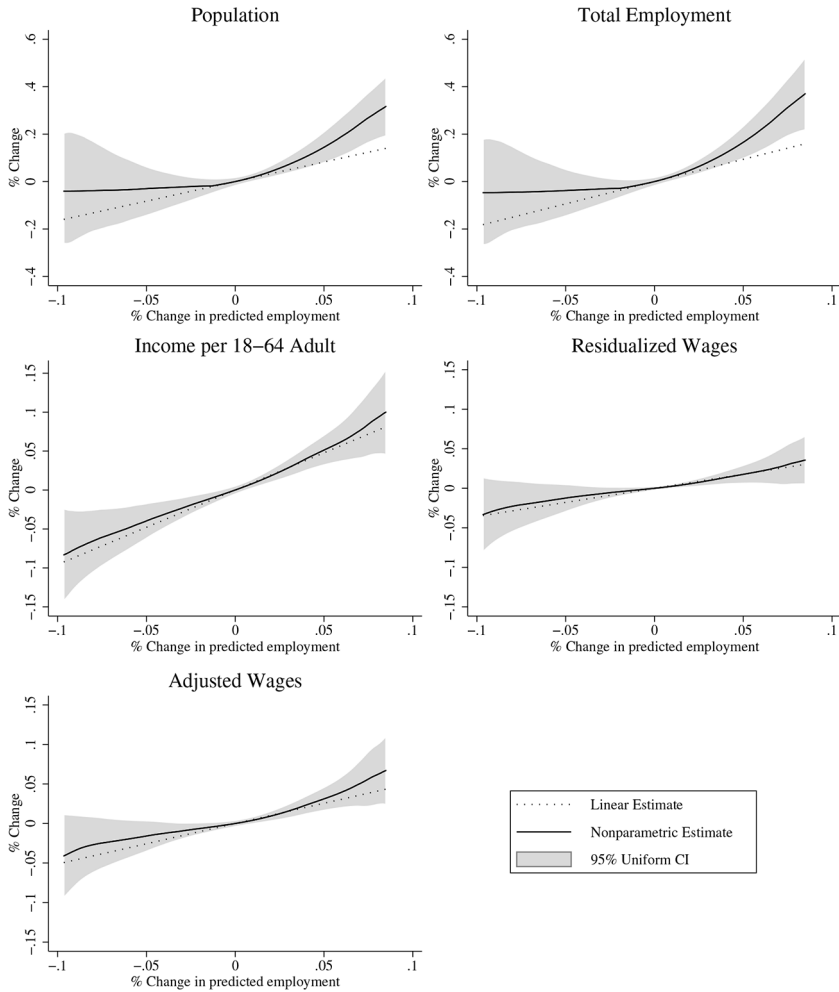


FIG. 4.—Reduced-form results. This figure reports nonparametric reduced-form estimates using US census data and Regional Economic Information System data. See the appendix (available online) for details on the data set. All graphs are nonparametric local linear regressions. All results include year fixed effects in the nonparametric model. The estimates are constrained to be monotonic following the rearrangement procedure of Chernozhukov, Fernandez-Val, and Galichon (2009). The 95% uniform confidence intervals are computed using 10,000 bootstrap replications, resampling metropolitan statistical areas with replacement. In each bootstrap step, an undersmoothed local linear bandwidth is chosen following Hall (1992).

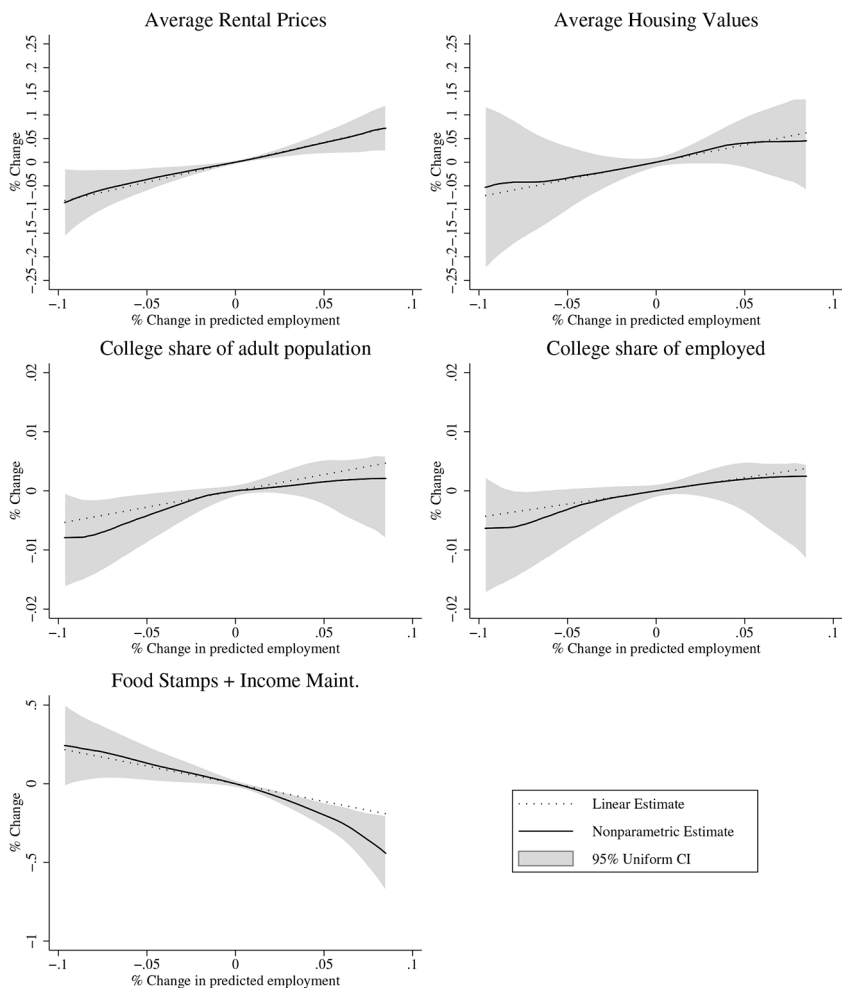


FIG. 5.—Reduced-form results, continued. This figure reports nonparametric reduced-form estimates using US census data and Regional Economic Information System data. See the appendix (available online) for details on the data set. All graphs are nonparametric local linear regressions. All results include year fixed effects in the nonparametric model. The estimates are constrained to be monotonic following the rearrangement procedure of Chernozhukov, Fernandez-Val, and Galichon (2009). The 95% uniform confidence intervals are computed using 10,000 bootstrap replications, resampling metropolitan statistical areas with replacement. In each bootstrap step, an undersmoothed local linear bandwidth is chosen following Hall (1992).

below will be useful in assessing whether the data reject the null hypothesis that the parametric linear model is appropriate. In all figures, the nonparametric estimates are local linear regressions, and the nonparametric estimates are constrained to be monotonic following the rearrangement procedure of Chernozhukov, Fernandez-Val, and Galichon (2009).²⁰

Overall, across all of the graphs the only clear evidence of an asymmetric response is for employment and population. The population and employment graphs show a convex relationship with the labor demand instrumental variable. By contrast, there is no evidence of a similar asymmetric relationship for housing values, rental prices, or any measure of wages (wage measures are defined below). As shown by the simulated data in figure 3, these results are consistent with a concave housing supply curve and limited mobility costs. To formally test for the existence of an asymmetric response (and measure the magnitude of the asymmetry when it exists), the next subsection reports results from quadratic specifications and nonparametric specification tests.

B. Reduced-Form Results

This section reports estimates of equation (7) to investigate the responsiveness of wages, employment, and population to changes in local labor demand. The baseline reduced-form estimating equation is reproduced here:

$$\Delta x_{it} = \beta \cdot \Delta \hat{\theta}_{it} + \delta \cdot (\Delta \hat{\theta}_{it})^2 + \alpha_t + \Delta \varepsilon_{it}.$$

The baseline results are reported in tables 2–4. Table 2 presents results for overall population, employment, and wages. Column 1 shows the results for the total population between the ages of 18 and 64.²¹ The estimate of β is precise and strongly statistically significant ($p < .001$), which verifies that the measure of predicted employment changes strongly predicts actual shifts in local population. The estimate of δ is also economically and statistically significant ($\hat{\delta} = 28.010$; standard error, 7.905). One way to interpret the magnitude of this estimate is to calculate the marginal effect at 1 standard deviation greater than zero and 1 standard deviation less than zero; these estimates are -0.152 and 3.757 , respectively, and the difference between these estimates is strongly statistically significant ($p < .001$).²² Additionally, a nonparametric specification test strongly rejects the null hypothesis that the relationship

²⁰ The rearranged estimates are more efficient under the null hypothesis that the true relationship is (weakly) monotonic. In general, the unconstrained nonparametric estimates are similar to the results reported in the figures in the main text.

²¹ Results using the population between the ages of 25 and 54 are very similar.

²² Note that the p -value for the test of whether the marginal effects are the same at 1 standard deviation above and below zero is exactly the same as the p -value for the test of whether the quadratic term is statistically significantly different from zero.

Table 2
Effects of Local Labor Demand Shocks

Dependent Variable	Population Employment (1) (2)	Employment-to- Population Ratio (3)	Income per 18-64- Year-Old Adult (4)	Residualized Aver- age Local Wage (5)	Residualized Wage, LFP Adjusted ("Adjusted Wage") (6)
% change in predicted employment (β)	1.825 (.442) [.000]	2.082 (.463) [.000]	.959 (.138) [.000]	.352 (.086) [.000]	.521 (.109) [.000]
(% change in predicted employment) ² (δ)	28.006 (7.901) [.000]	32.532 (9.098) [.000]	.378 (2.859) [.895]	-.757 (1.644) [.646]	1.455 (2.426) [.549]
(A) Marginal effect at $-\sigma$ (i.e., $\beta - 2\delta\sigma$)	-.130 (.842) [.878]	-.188 (.892) [.833]	.933 (.203) [.000]	.405 (.155) [.010]	.420 (.173) [.016]
(B) Marginal effect at $+\sigma$ (i.e., $\beta + 2\delta\sigma$)	3.779 (.539) [.000]	4.353 (.664) [.000]	.985 (.276) [.000]	.299 (.131) [.023]	.623 (.227) [.007]
p -value of test (A) = (B)	.000	.000	.895	.646	.549
p -value of nonparametric specification test	.000	.000	.492	.628	.451
R^2	.315	.354	.670	.471	.340
N	430	430	430	430	430

NOTE.—All columns report ordinary least squares results from estimating eq. (7). Data come from Integrated Public Use Microdata Series (IPUMS) 1980, 1990, and 2000 census extracts. The final sample is a balanced panel of 215 metropolitan statistical areas. The dependent variable is always the percentage change across periods except for col. 3, where it is the percentage point change. The residualized wage in col. 5 controls for observed compositional changes in the labor force between periods. The adjusted wage in col. 6 uses the residualized wage and additionally accounts for changes in labor force participation. The percentage change in predicted employment is formed by interacting cross-sectional differences in industrial composition with national changes in industry employment shares. The nonparametric specification test the null hypothesis that a linear model is appropriate against a nonparametric alternative. For more details, see the main text and sec. A.1 ("Data Appendix") of the appendix (available online). All specifications include year fixed effects. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each metropolitan area over time, are in parenthesis, and p -values are in brackets.

is linear in favor of a nonparametric alternative ($p < .001$).²³ In other words, the results in this column suggest that positive changes in local labor demand increase population more than negative changes reduce population. The results for employment in column 2 show evidence of a similar convex relationship. The results in column 3 using the percentage point change in the employment-to-population ratio show that not all of the reduction in local employment from an adverse shock comes from net out-migration; there is also a decline in labor force participation.

The remaining columns of table 2 explore the consequences of local labor demand shifts on wages. There are (at least) two difficulties in constructing an appropriate wage measure. The first difficulty is that the labor demand shock may induce compositional changes in the population, so that the change in the average wage will be confounded by composition effects. The second difficulty is that changes in labor force participation reduce income per adult but would be excluded using a measure of average wages based only on employed workers.

I approach these problems by first presenting two measures of changes in wage income that should represent upper and lower bounds of the true change in income holding characteristics of the workers fixed. The first measure (following Bound and Holzer [2000]) is the total wage income per 18–64-year-old adult. This measure will account for demand-induced changes in labor force participation but will also include compositional changes. The results are in column 4 and show a large effect of local labor demand on wages ($\hat{\beta} = 0.959$; standard error, 0.137). The second measure (following Shapiro [2006] and Albouy [2009, 2016]) uses the individual-level census data and regresses log wages of employed workers on a large set of controls and MSA fixed effects (for details, see app. sec. A.1 [“Data Appendix”]). The MSA fixed effect estimated from this regression is a composition-adjusted measure of the wage premium that I define as the “residualized wage.”²⁴ The results in column 5 using this measure show a much smaller wage response ($\hat{\beta} = 0.353$;

²³ I use the nonparametric specification test procedure suggested by Ellison and Ellison (2000), which groups the data into “bins” and creates a test statistic that is asymptotically distributed as a standard normal random variable. To my knowledge, there is not a data-driven procedure to select the proper bin width; therefore, I view the nonparametric specification test as complementary to the quadratic specification. While the nonparametric specification test does not rely on a specific parametric alternative, it is not possible to ensure that I have the right size and power in constructing my statistical tests. In almost all of the results that follow, inference based on the quadratic specification and the nonparametric specification test is similar.

²⁴ This measure is similar to the local wage premiums calculated in Shapiro (2006) and Albouy (2009, 2016). This measure does not control for unobservable changes in the composition of labor force. If unobservable changes in composition of labor force move in the same direction as observable changes, then the measured response of wages will be upward biased, and estimates of mobility costs will be conservative.

standard error, 0.086). However, this second measure does not account for changes in labor force participation. Assuming that at least some of the observed change in labor force participation is involuntary, then this measure will understate the total effect of the demand shock. To address this concern, I take the residualized wage measure and multiply it by the observed labor force participation rate.²⁵ I call this the “adjusted wage” and use this as the preferred wage measure. This measure accounts for both compositional changes in the labor force in response to the shock and changes in labor force participation, and it therefore essentially assumes that reservation wages are negligible. Consequently, I expect this measure to provide an overestimate of mobility costs when I ultimately estimate the full model via GMM. As a way of bounding the estimated magnitude of mobility costs, I also report GMM estimates that use the residualized wage instead of the adjusted wage. The residualized wage will give a lower bound on the estimated magnitude of mobility costs, as it assumes that reservation wages are approximately equal to accepted wages for all employed workers.

As expected, the magnitude of the effect of local labor demand for adjusted wages lies in between the other two wage measures ($\hat{\beta} = 0.520$; standard error, 0.109). Since the magnitude of changes in labor force participation is modest, the estimates for adjusted wages are closer to the estimates for residualized wages than the estimates using the per capita income measure. Regardless of the measure of wages used, however, the important conclusion that emerges from columns 4–6 is that there is no evidence of an asymmetric response of wages to shifts in local labor demand in any of the wage measures. It is only population and local employment that respond asymmetrically.

Table 3 reports results on population, employment, and wages separately for high-skill and low-skill workers. I define low-skill workers as those without a college degree and high-skill workers as those with at least a college degree. The patterns in table 2 are reproduced when looking separately within each skill group: population and employment respond asymmetrically, and there is no evidence of a similar asymmetric response for either high-skill or low-skill wages. Furthermore, the magnitude of the wage effects are similar across high-skill and low-skill workers, consistent with the assumption that the labor demand shifts are factor neutral.²⁶ Additionally, columns 5 and 6 show suggestive evidence that the skill composition of the adult population and labor force also responds asymmetrically. In other words, negative shocks reduce the college share of adult population more than positive

²⁵ Note that when I present results by skill below, I use the labor force participation rate in the given skill group to adjust the residualized wage measure.

²⁶ Results from stacked regressions do not reject the null hypothesis that the average wage response for high-skill workers is the same as the average wage response for low-skill workers ($p = .523$).

Table 3
Effects of Labor Demand Shocks by Skill

Dependent Variable	Adult Population		Total Employed		College Share		Residualized Wage		Adjusted Wage	
	College (1)	Noncollege (2)	College (3)	Noncollege (4)	Adult Population (5)	Employed Population (6)	College (7)	Noncollege (8)	College (9)	Noncollege (10)
% change in predicted employment (β)	1.953 (.543)	1.631 (.433)	2.226 (.552)	1.886 (.455)	.050 (.024)	.038 (.027)	.295 (.079)	.345 (.085)	.466 (.100)	.515 (.108)
(% change in predicted employment) ² (δ)	35.199 (10.361)	28.053 (7.688)	36.975 (10.990)	32.890 (8.831)	-.816 (.349)	-1.022 (.376)	-1.321 (1.399)	-.831 (1.628)	-.497 (2.012)	1.521 (2.371)
(A) Marginal effect at $-\sigma$	-.503 (.982)	-.327 (.838)	-.354 (.991)	-.409 (.885)	.107 (.033)	.110 (.037)	.388 (.139)	.403 (.152)	.501 (.167)	.409 (.168)
(B) Marginal effect at $+\sigma$	4.409 (.818)	3.589 (.498)	4.806 (.896)	4.181 (.624)	-.007 (.036)	-.033 (.038)	.203 (.111)	.287 (.130)	.432 (.177)	.622 (.223)
p -value of test (A) = (B)	.001	.000	.001	.000	.020	.007	.346	.610	.805	.522
p -value of nonparametric specification test	.000	.000	.000	.000	.134	.073	.621	.612	.605	.428
R^2	.558	.240	.569	.262	.772	.751	.432	.659	.472	.210
N	430	430	430	430	430	430	430	430	430	430

NOTE.—All columns report ordinary least squares results from estimating eq. (7). Data come from Integrated Public Use Microdata Series (IPUMS) 1980, 1990, and 2000 census extracts. The final sample is a balanced panel of 215 metropolitan statistical areas. The dependent variable is always the percentage change across periods except in cols. 5 and 6, which report percentage point changes in the college share. The percentage change in predicted employment is formed by interacting cross-sectional differences in industrial composition with national changes in industry employment shares. For more details, see table 2, the main text, and sec. A.1 (“Data Appendix”) of the appendix (available online). All specifications include year fixed effects. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each metropolitan area over time, are in parenthesis, and p -values are in brackets.

shocks increase college share. I emphasize that this asymmetric response is not as robust as the estimated asymmetric responses for population and employment for each skill group. As the simulations in figure 3 make clear, when an asymmetric population responses arises from either a concave housing supply curve or heterogeneous costs of out-migration, there is (at most) a small asymmetric responses in the high-skill population share.

Next, table 4 looks at three important nonlabor outcomes: real estate rental prices, housing values, and aggregate expenditures on public assistance programs. The measures of average rental prices and housing values are purged of observable changes in the quality of the housing stock following a procedure similar to the one used to create the residualized wage measure (for details, see app. sec. A.1 ["Data Appendix"]). Column 1 in table 4 reports

Table 4
Effects of Labor Demand Shocks on Housing Market
and Public Assistance Expenditures

Dependent Variable	Residualized Rental Prices (1)	Residualized Housing Values (2)	Food Stamps + Income Maintenance Expenditures (3)
% change in predicted employment (β)	.842 (.151) [.000]	.712 (.360) [.049]	-2.384 (.616) [.000]
(% change in predicted employment) ² (δ)	-1.006 (2.758) [.716]	-2.759 (6.314) [.663]	-21.775 (12.138) [.074]
(A) Marginal effect at $-\sigma$.912 (.241) [.000]	.905 (.576) [.118]	-.865 (1.022) [.398]
(B) Marginal effect at $+\sigma$.771 (.248) [.002]	.520 (.562) [.356]	-3.904 (1.072) [.000]
p -value of test (A) = (B)	.716	.663	.074
p -value of nonparametric specification test	.596	.295	.241
R^2	.099	.144	.403
N	430	430	430

NOTE.—All columns report ordinary least squares results from estimating eq. (7). Data come from Integrated Public Use Microdata Series (IPUMS) 1980, 1990, and 2000 census extracts and the Regional Economic Information System (REIS) database. The REIS database contains total county-level expenditures on food stamps and income maintenance programs. These data are aggregated to metropolitan statistical areas (MSAs) using the 1990 MSA definition and are adjusted per non-college-educated capita using MSA population estimates from the census. The final sample is a balanced panel of 215 MSAs. The dependent variable is always the percentage change across periods. The percentage change in predicted employment is formed by interacting cross-sectional differences in industrial composition with national changes in industry employment shares. For more details, see table 2, the main text, and sec. A.1 ("Data Appendix") of the appendix (available online). All specifications include year fixed effects. Standard errors, adjusted to allow for an arbitrary variance-covariance matrix for each metropolitan area over time, are in parenthesis, and p -values are in brackets.

results for rental prices, which respond strongly to local labor demand. The results for housing values in column 2 are similar in magnitude although somewhat less precise. As with the wage results, there is no evidence of an asymmetric response in either column; the estimates of δ are statistically insignificant and at most modest in magnitude, and the nonparametric specification tests fail to reject the parametric (linear) model in both columns.²⁷ Table A5 (tables A1–A8 are available online) reports similar results using the unconditional average rental prices and average housing values, as well as results using the repeated-sales housing price index published by the Federal Housing Finance Agency, formerly the Office of Federal Housing Enterprise Oversight. Consistent with the results in table 4, there is no evidence of an asymmetric response in any of these alternative specifications.

Last, column 3 reports estimates using aggregate expenditures on food stamps and income maintenance programs. The results show that expenditures on these programs respond strongly to local labor market conditions. The estimated magnitude of the response is large ($\hat{\beta} = -2.367$) and implies that a 1% decline in local labor demand increases aggregate expenditures on these two programs by 2.4%. Although the quadratic term is marginally significant ($p = .074$), the nonparametric test does not reject the linear model ($p = .241$), suggesting that the nonlinear relationship estimated in the quadratic specification is not robust.²⁸

A setting in which population and employment respond asymmetrically to positive and negative labor demand shocks while wages, rental prices, and housing values respond symmetrically is consistent with the model simulation where mobility costs are limited and the housing supply curve is concave. Before moving beyond this qualitative conclusion to quantitative estimates of mobility costs and housing supply curve parameters, I next document that these reduced-form results are not driven by unobserved trends, outliers, sample selection, or heterogeneous industry-specific effects. After that, I conclude by estimating the full model described above using a nonlinear GMM estimator.

C. Robustness

This section summarizes the results of additional analysis that assesses the robustness of the main results; the results are presented in the appendix.

²⁷ Additionally, results from stacked regressions reject that the quadratic terms are the same for population and rental prices ($p = .0004$) and reject that the quadratic terms are the same for population and housing values ($p = .001$).

²⁸ Table A6 reports estimates for various other transfer programs, including Medicare, disability benefits, SSI, and veterans benefits, and the results are qualitatively similar. I focus on food stamps and income maintenance income because these programs are explicitly designed to smooth consumption.

1. Industry Trends

First, I categorize industries on the basis of their decadal changes in total national employment, and I group industries into four categories: persistently expanding/declining industries, stable industries, volatile industries, and other industries.²⁹

The top 20 industries according to average national employment share in each of these categories are listed in table A1. The industries in each of the categories match expectations given the secular industry trends during this time period. Persistently expanding industries are concentrated in services, health care, data processing, and leisure goods, while persistently contracting industries are in apparel, publishing, manufacturing, and tobacco. Volatile industries include natural resource industries such as oil and gas extraction as well as defense industries.

I next construct predicted employment by excluding the persistently expanding/declining industries.³⁰ The resulting relative demand index is then purged of any variation caused by secular trends in health care, services, and manufacturing. Table A2 reports results from estimating equation (7) using this alternative measure of predicted employment as an instrumental variable for local labor demand. Overall, the results are fairly similar using this alternative instrumental variable. Moreover, the correlation between the labor demand instrument used in columns 2 and 5 is 0.48, suggesting that the similarity across columns is not simply a mechanical consequence of the different instruments exploiting similar sources of variation.

Last, table A3 reports similar results using an alternative construction of the instrumental variable that replaces industry employment (and employment shares) with industry payroll. This construction follows Katz and Murphy (1992), which measures labor demand shocks using changes in industry payroll (rather than employment). I find that the same asymmetries in employment and population show up using this alternative instrumental

²⁹ The categories are defined as follows: (1) persistently expanding/declining industries are industries where employment either increased in every decade or decreased in every decade; (2) stable industries are industries where employment did not increase or decrease more than 20% in any of the decades, whether or not they are persistently expanding or declining; (3) volatile industries are industries that experienced employment growth of more than 20% and decreases of more than 20% during the sample period; and (4) other industries are industries not otherwise categorized.

³⁰ Formally, predicted employment growth is computed by using only the subset of industries that pass a given filter:

$$\pi'_{i,t} = \sum_{k \in K' \subset K} \varphi_{i,k,t-\tau} \left(\frac{v_{-i,k,t} - v_{-i,k,t-\tau}}{v_{-i,k,t-\tau}} \right),$$

where K is the set of all industries and K' is the set of industries that pass the filter.

variable, and I also find no evidence of similar asymmetries for income per capita, average wage, and average rental prices.

2. *Alternative Specifications*

I next turn to reporting alternative specifications that vary the sample definition and the set of time-varying controls. Table A4 reports results that expand sample (by adding data on 2000–2007 change and adding “psuedo-MSAs” that group together all individuals in a state who are not living in an identifiable MSA). Additionally, this table results from alternative control variables, including region-by-year fixed effects and MSA-specific linear time trends. This table also reports results using the County Business Patterns (CBP) data set to construct the local labor demand instrument rather than using census data (for details, see app. sec. A.1 [“Data Appendix”]). The CBP data contain finer industry categories, which in principle could reduce measurement error, but there are two primary drawbacks: a high rate of suppressed data at the county-by-industry level and the county-level data must be aggregated to the MSA level. Last, table A7 reports results that drop each one (of nine) census regions.

In both tables, the reduced-form results consistently show a pattern of a significant asymmetric response of population and employment to changes in local labor demand. This robust result contrasts with a lack of similar asymmetric responses for wages, housing values, and rental prices.

V. GMM Estimates

The reduced-form results presented above directly test for the existence of asymmetric responses of wages, population, employment, and housing prices to symmetric labor demand shocks. These results do not estimate any of the economic parameters in the theoretical model and are therefore not quantitatively informative about the distribution of mobility costs by skill and the actual incidence of labor demand shocks. This section reports results from a joint estimation of the full model using a nonlinear simultaneous-equations GMM estimator. The econometric setup follows from the theoretical model presented above and imposes moment conditions that can be used to identify the parameters of interest. In particular, the GMM estimator can recover flexible estimates of the housing supply curve and mobility cost functions for high-skill and low-skill workers. These estimates can be used to assess the relative importance of housing expenditures, transfer payments, and mobility costs in generating the observed migration patterns in the data. Additionally, because I parameterize the model so that there are more moment conditions than (remaining) parameters to estimate, the GMM estimator admits a χ^2 overidentification test of the full model.

To implement the GMM estimator, the following equations (derived from eqq. [1]–[6] in the model above) are used:

$$\begin{aligned}
\Delta e_{it}^{wH} &= \Delta w_{it}^H - (\Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)(\pi))\Delta H_{it} \\
&\quad + (\alpha - \rho)(1 - \pi)\Delta L_{it}), \\
\Delta e_{it}^{wL} &= \Delta w_{it}^L - (\Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)(1 - \pi))\Delta L_{it} \\
&\quad + (\alpha - \rho)(\pi)\Delta H_{it}), \\
\Delta e_{it}^b &= \Delta b_{it} - \Psi \Delta w_{it}^L, \\
\Delta e_{it}^H &= \Delta p_{it}^H + \Delta H^s(\Delta p_{it}^H) - (\nu(\Delta w_{it}^H + \Delta H_{it}) \\
&\quad + (1 - \nu)((1 - s_b^L)\Delta w_{it}^L + s_b^L \Delta b_{it} + \Delta L_{it})), \\
\Delta e_{it}^H &= \Delta w_{it}^H - s_{it}^H \Delta p_{it}^H + c^H(\Delta H_{it}), \\
\Delta e_{it}^L &= (1 - s_b^L)\Delta w_{it}^L + s_b^L \Delta b_{it}^L - s_{it}^L \Delta p_{it}^H + c^L(\Delta L_{it}),
\end{aligned}$$

where i indexes cities, t indexes time, and Δe_{it}^j represent error terms uncorrelated with shifts in labor demand.³¹ These equations jointly solve the local general equilibrium problem of how wages, employment, housing prices, and transfer payments respond to an exogenous labor demand shift $\Delta \theta_{it}$. The six endogenous variables are the following: Δp_{it}^H , Δw_{it}^H , Δw_{it}^L , ΔH_{it} , ΔL_{it} , and Δb_{it} . Note that the error terms are allowed to be freely correlated with each other, which gives rise to simultaneity bias that the GMM estimator is intended to address. The unknowns in the model are the following parameters and functions:

- Transfer income and housing expenditure shares (s_b^L , s_{it}^L , s_{it}^H)
- Aggregate share parameters (μ , ν)
- Labor demand parameters (α , ρ , π , ζ)
- Transfer payment elasticity (Ψ)
- Mobility cost functions ($c^L(\cdot)$ and $c^H(\cdot)$)
- Housing supply function ($\Delta H^s(\cdot)$)

To reduce the number of parameters to estimate, I calculate values of s_b^L , s_{it}^L , and s_{it}^H on the basis of external information. I compute $s_b^L = 0.05$ by dividing aggregate expenditures on food stamps and income maintenance programs by the sum of these expenditures and aggregate low-skill wage income. For

³¹ Each of these equations can be derived formally by including error terms that proportionally shift production, housing demand, housing supply, transfer payments, and indirect utility. For example, redefine the equilibrium condition for transfer payments as follows: $b_{it}^L = e_{it}^b \cdot B^L(w_{it}^L)^{\Psi^L}$, where e_{it}^b is a random variable that represents unobservable shocks to transfer payment expenditures (and $E[e^i] = 1$). Totally differentiating this condition gives the expression $\Delta b_{it}^L = \Psi^L(\Delta w_{it}^L) + \Delta e_{it}^b$, which is the equation used in the GMM estimation.

the housing expenditure shares, I use $s_{H_i}^L = 0.34$ for non-college-education households and $s_{H_i}^H = 0.30$ for college-educated households.³²

For the labor demand curve, I compute $\pi = 0.37$ on the basis of average wages for high-skill and low-skill workers and average share of high-skill workers in the adult population. I compute the wage premium (ζ) as 1.75, which is the average wages of college-educated workers divided by the average wages of non-college-educated workers. I next compute the average share (over this time period) of college-educated workers in the labor force (μ) as 0.25. Using the formula for π in section II, this gives $\pi = 0.37$. I compute $\nu = 0.34$ on the basis of the average wages, the skill share, and the housing expenditure shares from above. I choose $\rho = 0.29$ on the basis of Katz and Murphy (1992).³³ This leaves the returns-to-scale parameter (α) to be estimated. Although this parameter will be estimated from functional form assumptions, it is still useful to include the two moments of the labor demand curve to check the overall fit of the model.³⁴ This means that misspecification in the functional form of labor demand equation will cause bias estimates in all of the parameters when estimating the entire system of equations. Therefore, I also report results below that drop the labor demand moments.³⁵

Finally, I choose the following functional forms for the mobility cost functions and housing supply elasticity:

$$c^j(x) = \frac{\sigma^j(\exp(\beta^j x) - 1)}{\beta^j} \quad j \in \{L, H\},$$

$$\Delta \mathcal{H}^s(x) = \frac{\sigma^H(\exp(\beta^H x) - 1)}{\beta^H}.$$

³² Average household income is \$82,439 for high-skill households in the baseline sample and is \$48,456 for low-skill households. Assuming $s_{H_i}^H = 0.30$ for high-skill households and an income elasticity of 0.8, then $s_{H_i}^L = 0.34$ for low-skill households. This is similar to differences in housing expenditure shares measured in the Consumer Expenditure Survey data comparing households based on the highest level of education of any member of the household.

³³ Katz and Murphy (1992) estimate the elasticity of substitution between high-skill and low-skill labor ($\sigma_{H,L}$) to be 1.4. This gives $\rho = 1 - 1/\sigma_{H,L} = 0.29$.

³⁴ Since the instrumental variable shifts the labor demand curve, parameters of the labor demand curve itself are identified from functional form assumptions.

³⁵ Because the labor demand instrument is measured with error, when using it in the GMM estimation I rescale it by regressing adjusted wages on the instrument and scale the instrument so that this regression with the rescaled instrument would give a coefficient of 1.0. An alternative is to modify the labor demand moments to include an additional parameter (κ) as follows:

$$\Delta e_{it}^{wH} = \Delta w_{it}^H - (\kappa \Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)(\pi)) \Delta H_{it} + (\alpha - \rho)(1 - \pi) \Delta L_{it}),$$

$$\Delta e_{it}^{wL} = \Delta w_{it}^L - (\kappa \Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)(1 - \pi)) \Delta L_{it} + (\alpha - \rho)(\pi) \Delta H_{it}).$$

This procedure yields very similar results.

These functions are the exponential transformations suggested by Manly (1976), which represent Box-Cox transformations of exponentiated variables and are defined so that if $\beta^j = 0$, then the functions simplify to $\sigma^j x$. These functions are flexible enough to accommodate interesting curvature with only two parameters, and they are everywhere monotonic and have continuous first derivatives, which greatly simplifies the computation.

Ultimately, there are eight remaining parameters to estimate $\{\sigma^H, B^H, \sigma^L, \beta^L, \sigma^H, \beta^H, \Psi, \alpha\}$: two housing supply curve parameters (σ^H, B^H), two low-skill mobility cost parameters (σ^L, β^L), two high-skill mobility cost parameters (σ^H, β^H), the responsiveness of transfer payments to low-skill wages (Ψ), and the returns-to-scale parameter (α). The resulting GMM estimator solves a nonlinear simultaneous-equations problem, so to estimate the nonlinear parameters I take nonlinear functions of the instrumental variable ($\Delta\theta$) to achieve identification. I use $\Delta\theta, (\Delta\theta)^2, (\Delta\theta)^3, (\Delta\theta)^4$, and $(\Delta\theta)^5$ as instrumental variables.³⁶ This results in 30 moment conditions (the five polynomial functions of the instrument times the six error terms). The full model is estimated using a standard two-step GMM procedure (see sec. A.5 of the appendix for details of this procedure).

The GMM estimates are presented in table A8. The main results suggest that the local housing supply curve is concave ($\beta^H = 6.306$; standard error, 1.774). One way to interpret the housing supply coefficients is to compute the increase in housing supply when housing prices exogenously rise by 20% (24.1%) and compare it to the decrease in housing supply when housing prices decline by 20% (-6.8%). In other words, the magnitude of housing supply response is about four times larger for an increase in housing prices than for an equal-sized decrease in housing demand.

The estimates of the mobility cost function parameters show no evidence of an asymmetric mobility cost function for either high-skill or low-skill workers; the estimates suggest that the mobility cost functions are approximately linear. The point estimates for σ^L and σ^H are precisely estimated and statistically significantly different from zero, suggesting the existence of nonnegligible mobility costs. To get a sense of the magnitudes, the point estimates imply that the 10th percentile of mobility costs in a city (i.e., the marginal migrant after 10% of the population has out-migrated following a negative shock) is roughly 17.4% of annual income for high-skill workers and 17.0% of annual income for low-skill workers. In other words, despite the fact that low-skill workers are disproportionately likely to remain in declining cities following negative shocks, the point estimates imply that high-skill workers have very similar mobility costs as a fraction of income (and therefore that low-skill workers have

³⁶ In principle, only the quadratic term is needed to identify the parameters of the model. Adding additional polynomial terms increases statistical power at the cost of introducing bias (either because the orthogonality assumption is not satisfied at higher moments or the additional polynomial introduces a weak-instruments problem).

lower absolute mobility costs on average). Next, I estimate a large and statistically significant transfer payment elasticity parameter; the coefficient implies that a 1% decline in low-skill wages increases transfer payment expenditures by 3.8% (standard error, 0.5%). Last, the overidentification test does not reject the null hypothesis that the deviations of the empirical moments from the model are due to chance ($p = .515$).³⁷

Table A8 also reports GMM estimates under alternative assumptions, such as different (counterfactual) assumptions about housing expenditure shares, transfer payments, and mobility costs. The results are summarized in figures 6 and 7. The estimated mobility cost functions in these different scenarios are graphed in figure 6. Both the figure and the model estimates suggest that transfer payments are responsible for a majority of the relative difference in mobility by skill. However, the magnitudes of mobility cost estimates are much larger for both types of workers when housing expenditures are ignored. In other words, the asymmetric population response for both high-skill and low-skill workers is primarily due to the asymmetric housing supply curve, while the differential response by skill is primarily due to transfer payments.³⁸

The final rows of table A8 explore different assumptions about the elasticity of substitution between high-skill and low-skill workers and also explores using an alternative measure of wages. I also explore not using any of the labor demand curve moments. All of these alternative specifications lead to fairly similar estimates of housing supply curve and mobility costs.³⁹

³⁷ The model estimates are also not able to reject constant returns to scale ($p = .129$). This is consistent with the reduced-form results, which found no evidence of an asymmetric response of wages. Average wages did not respond asymmetrically, but population and employment did, which is consistent with constant returns to scale. If there were decreasing returns to scale, then the asymmetric response of employment to the local labor demand shock would imply an asymmetric wage response as well.

³⁸ I also experiment with assuming the demand for housing is homothetic, so that the housing expenditure shares are the same across the two skill groups. I choose $s_H^L = s_H^H = 0.33$ to match the average housing expenditure share across the entire population, and the results are fairly similar to the baseline results in the first row, implying that the nonhomotheticity assumed in the baseline model does not substantially account for the differential out-migration rates by skill.

³⁹ Interestingly, the fit of the model is best when using $\sigma_{H,L} = 1.4$ (row 1) following Katz and Murphy (1992), as opposed using either of the two extreme values of $\sigma_{H,L}$. The final row of table A8 reports estimates that drop the labor demand curve moments. The reason why alternative assumptions regarding the elasticity of substitution do not substantially affect the estimated mobility cost functions is that the labor demand moments contribute only indirectly through the optimal GMM weighting matrix estimated in the first step of the two-step procedure. Thus, it is not surprising that dropping the labor demand moments entirely does not significantly affect the estimates of the mobility cost functions.




FIG. 6.—Generalized method of moments (GMM) estimates of the full model. *A* presents the housing supply curve that is estimated in the baseline model (table A8, row 1). *B* and *C* report estimated mobility functions under various assumptions about housing expenditure shares and transfer payments. See section V and the appendix (available online) for more details on the GMM estimation.

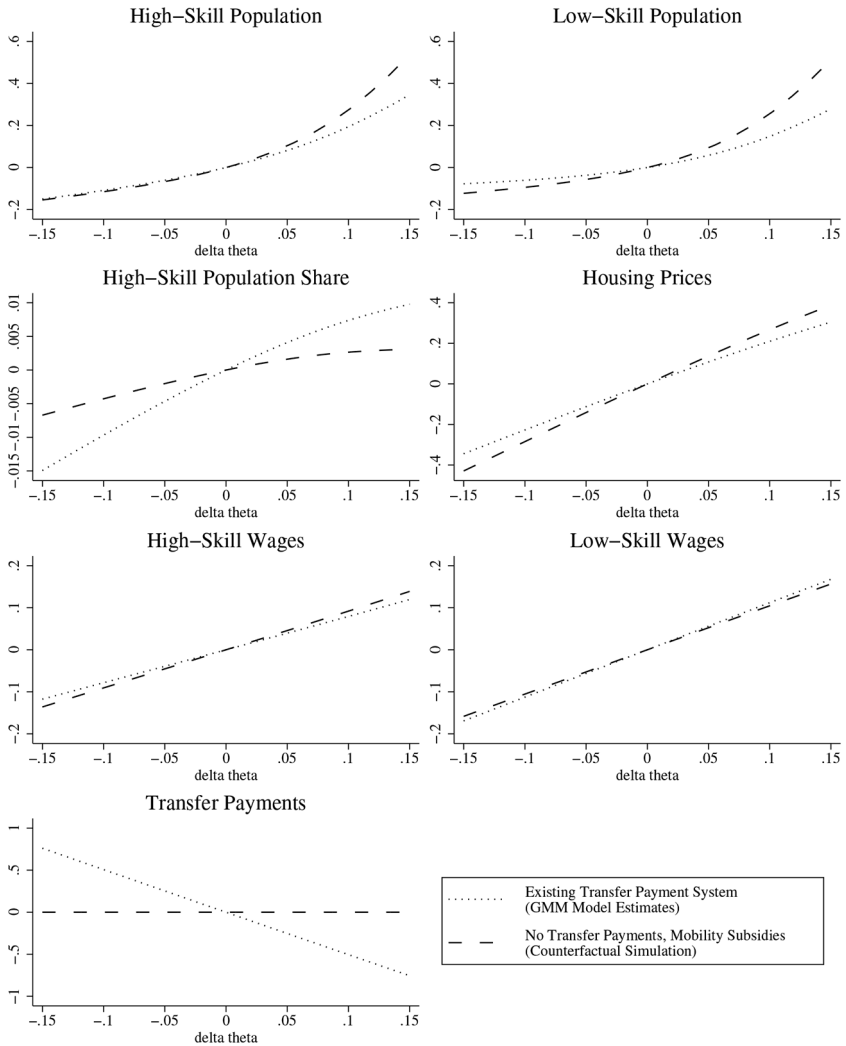


FIG. 7.—Counterfactual simulations. This figure reports simulations based on generalized method of moments (GMM) estimates of the full model. The GMM estimates are used to run simulations similar to those presented in figure 3. The graphs report results of two simulations: (1) simulation based on estimates of the baseline GMM model using the existing transfer payment system and (2) counterfactual simulation based on the same estimates but the transfer payment system is replaced with mobility subsidies that reduce mobility costs by 50%.

One use of the GMM estimates is to construct out-of-sample counterfactual simulations of alternative policies regarding social transfers. Figure 7 reports results from one such simulation. In this simulation, the system of means-tested transfers (summarized by the parameter Ψ) has been replaced by a system of mobility subsidies that reduces the mobility costs of all workers by 50%.⁴⁰ Each panel in the figure shows the response of a different endogenous variable. The figure shows that the mobility subsidies increase the magnitude of low-skill out-migration following adverse shocks relative to the system of means-tested transfer payments. Therefore, the high-skill population share is much less responsive to shifts in local labor demand with mobility subsidies. One motivation for such a policy would be if there exist strong negative externalities from increasing the concentration of low-skill workers in a particular area; in this case, mobility subsidies appear to provide wage insurance to low-skill workers without differentially reducing their incentive to out-migrate.

VI. Conclusion

Low-skill workers are comparatively immobile. When labor demand slumps in a city, college-educated workers tend to relocate, whereas non-college-educated workers are disproportionately likely to remain to face declining wages and employment. These facts may indicate that mobility is disproportionately costly for low-skill workers. This paper proposes and tests an alternative explanation, which is that the incidence of adverse labor demand shocks is borne in large part by (falling) real estate rental prices and (rising) social transfers. The spatial equilibrium model developed in this paper illustrates how wages, employment, population, housing prices, and transfer payments reequilibrate after a local labor demand shock. Appropriately parameterized, the model identifies both the magnitude of unobserved mobility costs by skill and the shape of the local housing supply curve. The model estimates suggest that the main explanation for the comparative immobility of low-skilled workers is not higher mobility costs per se but rather a lower incidence of adverse local demand shocks.

The reduced-form analysis in this paper documents asymmetric responses to local labor demand shocks and uses these responses to learn about the incidence of local labor demand shocks. Subsequent to the circulation of this paper, broadly similar asymmetric responses have also been documented by Dao, Furceri, and Loungani (2017) using a reduced-form vector autoregression model and a similar instrumental variables strategy. Taken together, the findings in Dao, Furceri, and Loungani (2017) and this paper suggest that

⁴⁰ Although this is an obviously stylized form of mobility subsidies, it is a rough approximation of a policy that took the form of a tax credit that was indexed to income. Recall that mobility costs in the model are defined as a fraction of annual income.

population gains in local markets experiencing positive labor demand shocks are generally larger in magnitude than population losses in areas experiencing negative shocks. These asymmetries are consistent with the model of durable housing in Glaeser and Gyourko (2005) and the asymmetric housing supply curve estimated in this paper.

I conclude with a discussion of some possible directions for future work. One area involves studying individual transfer programs in more detail. For example, take-up of the Social Security disability insurance (SSDI) program is usually an “absorbing state,” as once individuals receive SSDI, they rarely leave the program (Autor and Duggan 2006). Based on this understanding, the econometric approach in this paper could be used to test whether positive shifts in local labor demand affect SSDI take-up less than negative shifts.

Another important area of future work involves extending the spatial equilibrium model in this paper. Some work in this area has already taken place subsequent to the circulation of this paper. In particular, several recent papers have constructed richer and more sophisticated spatial equilibrium models that incorporate endogenous local amenities (Diamond 2016), imperfect mobility of firms as well as workers (Suárez Serrato and Zidar 2016), and state and local taxes (Fajgelbaum et al. 2019). Each of these papers follows this paper in taking a GMM approach to the estimation of the spatial equilibrium model parameters, but each of these papers provides micro foundations of individual behavior, which enable these papers to consider richer counterfactual scenarios.

There are other extensions to the spatial equilibrium model that may also be worth exploring. The model in this paper does not distinguish between homeowners and renters, which may rule out local adjustments to shocks that differ between these groups. On the one hand, homeowners’ “user cost” of housing has declined following a negative labor demand shock, which is similar to the decline in rental prices experienced by renters. However, unlike renters, homeowners experience a negative wealth effect from a decline in housing values, which may affect their responsiveness to local labor demand shocks. The model in this paper also does not allow for extensive margin labor supply adjustments. This is a natural extension to consider, since the reduced-form empirical results suggest that negative local labor demand shocks appear to reduce wages, employment, and labor force participation. Incorporating each of these additional features may help deliver a more complete understanding of the incidence of local labor demand shocks.

References

- Albouy, David. 2009. The unequal geographic burden of federal taxation. *Journal of Political Economy* 117, no. 4:635–67.
- . 2016. What are cities worth? Land rents, local productivity, and the total value of amenities. *Review of Economics and Statistics* 98, no. 3:477–87.

- Artuç, Erhan, Shubham Chaudhuri, and John McLaren. 2010. Trade shocks and labor adjustment: A structural empirical approach. *American Economic Review* 100, no. 3:1008–45.
- Autor, David, and Mark Duggan. 2006. The rise in disability rolls and the decline in unemployment. *Quarterly Journal of Economics* 118, no. 1:157–205.
- Bayer, Patrick, Nathaniel Keohane, and Christopher Timmins. 2008. Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management* 58, no. 1:1–14.
- Bartik, Timothy J. 1991. Who benefits from state and local economic development policies. Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.
- Blanchard, Olivier Jean, and Larry Katz. 1992. Regional evolutions. *Brookings Papers on Economic Activity* 1992, no. 1:1–75.
- Bound, John, and Harry J. Holzer. 2000. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics* 18, no. 1:20–54.
- Busso, Matias, Jesse Gregory, and Patrick Kline. 2013. Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review* 103, no. 2:897–947.
- Chaudhuri, Shubham, and John McLaren. 2007. Some simple analytics of trade and labor mobility.” NBER Working Paper no. 13464, National Bureau of Economic Research, Cambridge, MA.
- Chernozhukov, Victor, Ivan Fernandez-Val, and Alfred Galichon. 2009. Improving estimators of monotone functions by rearrangement. *Biometrika* 96, no. 3:559–75.
- Dao, Mai, Davide Furceri, and Prakash Loungani. 2017. Regional labor market adjustment in the United States: Trend and cycle. *Review of Economics and Statistics* 99, no. 2:243–57.
- Diamond, Rebecca. 2016. The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000. *American Economic Review* 106, no. 3:479–524.
- Ellison, Glenn, and Sara Fisher Ellison. 2000. A simple framework for non-parametric specification testing. *Journal of Econometrics* 96, no. 1:1–23.
- Fajgelbaum, Pablo, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. 2019. Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms. *Review of Economic Studies* 86, no. 1:333–76.
- Glaeser, Edward L., and Joseph Gyourko. 2005. Urban decline and durable housing. *Journal of Political Economy* 113, no. 2:345–75.
- Hall, Peter. 1992. Effect of bias estimation on coverage accuracy of bootstrap confidence intervals for a probability density. *Annals of Statistics* 20, no. 2:675–94.

- Katz, Larry, and Kevin M. Murphy. 1992. Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics* 102, no. 1:35–78.
- Kennan, John, and James R. Walker. 2011. The effect of expected income on individual migration decisions. *Econometrica* 79, no. 1:211–51.
- . 2010. Wages, welfare benefits and migration. *Journal of Econometrics* 156, no. 1:229–38.
- Luttmer, Erzo. 2005. Neighbors as negatives: Relative earnings and well-being. *Quarterly Journal of Economics* 120, no. 3:963–1002.
- Malamud, Ofer, and Abigail Wozniak. 2012. The impact of college on migration: Evidence from the Vietnam generation. *Journal of Human Resources* 47:913–50.
- Manly, B. F. J. 1976. Exponential data transformations. *Journal of the Royal Statistical Society, Series D* 25, no. 1:37–42.
- Meyer, Bruce, Wallace K. C. Mok, and James X. Sullivan. 2009. The under-reporting of transfers in household surveys: Its nature and consequences. Harris School Working Paper no. 09.03.
- Moretti, Enrico. 2013. Real wage inequality. *American Economic Journal: Applied Economics* 5, no. 1:65–103.
- Roback, Jennifer. 1982. Wages, rents, and the quality of life. *Journal of Political Economy* 90, no. 6:1257–78.
- Ruggles, Steven, Matthew Sobek, Trent Alexander, Catherine A. Fitch, Ronald Goeken, Patricia Kelly Hall, Miriam King, and Chad Ronnander. 2004. Integrated Public Use Microdata Series: version 3.0 Machine-readable database. Minneapolis: Minnesota Population Center.
- Shapiro, Jesse. 2006. Quality of life, productivity, and the growth effects of human capital. *Review of Economics and Statistics* 88, no. 2:324–35.
- Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms. *American Economic Review* 106, no. 9:2582–624.
- Topel, Robert H. 1986. Local labor markets. *Journal of Political Economy* 94, no. 3:S111–S143.
- Wozniak, Abigail. 2010. Are college graduates more responsive to distant labor market opportunities? *Journal of Human Resources* 45, no. 4:944–70.