Available online at www.sciencedirect.com
ScienceDirect
Journal of Urban Economics 64 (2008) 178–195
JOURNAL OF Urban Economics
www.elsevier.com/locate/jue

Job creation and housing construction: Constraints on metropolitan area employment growth ☆

Raven E. Saks

Federal Reserve Board of Governors, 20th & C Street, NW, Washington, DC, USA

Received 8 August 2006; revised 17 December 2007

Available online 23 December 2007

Abstract

Differences in the supply of housing generate substantial variation in house prices across the United States. Because house prices influence migration, the elasticity of housing supply also has an important impact on local labor markets. I assemble evidence on housing supply regulations and examine their effect on metropolitan area housing and labor market dynamics. Locations with relatively few barriers to construction experience more residential construction and smaller increases in house prices in response to an increase in housing demand. Furthermore, housing supply constraints alter local employment and wage dynamics in locations where the degree of regulation is most severe.
Published by Elsevier Inc.

JEL classification: R23; R31; R52; J61

Keywords: Housing supply; Zoning; Local labor markets

1. Introduction

A growing literature argues that labor migration is one of the primary mechanisms through which metropolitan areas adjust to changes in local economic conditions (Blanchard and Katz, 1992; Gallin, 2004; Topel, 1986). Prospective migrants choose a location by comparing the benefit to living in each area to the cost of moving. Because housing is a large share of the household budget, house prices have an important effect on the relative value of wages across geographic areas.[1] As a result, areas with high house prices will attract fewer migrants holding other factors constant (Gabriel et al., 1993; Johnes and Hyclak, 1999).

Because housing markets influence migration, local employment growth depends critically on the capacity of the construction industry to accommodate increases in housing demand. In places where residential construction responds to new demand without difficulty, workers will move into the area with little change in house prices. In contrast, if new construction is con-

☆ The views in this paper are not necessarily those of the Board of Governors of the Federal Reserve or its staff.

E-mail address: raven.e.saks@frb.gov.

[1] In 2002, households in the United States spent 19.2% of average annual expenditures on shelter (US Bureau of Labor Statistics, 2002, Table 4).

0094-1190/$ – see front matter Published by Elsevier Inc.
doi:10.1016/j.jue.2007.12.003

strained, an increase in demand will lead mostly to higher house prices, with little change in employment. Thus, the elasticity of housing supply is a key factor in determining how labor markets adjust to changes in local economic conditions.[2]

Although changes in the housing supply are not the only way in which an area can adjust to a change in local labor demand, the correlation between changes in employment and new construction is strong. Controlling for year and metropolitan area fixed effects, a simple OLS regression of annual log changes in employment on annual log changes in the housing stock yields a coefficient of 0.57 (with a standard error of 0.03).[3,4] Therefore, growing cities must confront the issue of where new workers will live.

In this paper, I explore the effect of the housing supply on metropolitan area labor markets. To determine the elasticity of housing supply in individual metropolitan areas, I assemble evidence from six different sources of information on local land use policy. This new index reveals considerable heterogeneity in the extent of regulation across locations, and I find that areas with a larger degree of regulation experienced less residential construction and larger house price increases than less regulated locations from 1980 to 2000. In addition to these effects on local housing markets, I develop a simple model to show how the elasticity of housing supply (and consequently the degree of housing supply regulation) should also impact local labor markets. Specifically, a labor demand shock should result in lower employment growth, higher wages and higher house prices in places with an inelastic housing supply. Consistent with this theory, the long-run response of employment to an increase in labor demand is about 20 percent lower in metropolitan areas with a high degree of housing supply regulation.

2. Housing supply regulations across metropolitan areas

In order to explore the effect of the housing supply on labor markets, first it is necessary to identify credible variation in the elasticity of housing supply across locations. Using this variation, I will then show how the dynamics of metropolitan housing and labor markets depend on the responsiveness of the housing supply to shocks to labor demand. Empirical evidence on differences in the elasticity of housing supply across metropolitan areas is scarce because this parameter is not easy to observe.[5] Rather than estimating this elasticity from fluctuations in house prices and quantities, I use information on the restrictiveness of land use regulation in each location to evaluate the responsiveness of the housing supply in individual metropolitan areas. In this section, I describe how I construct this index and show that locations with more regulation have higher house prices and less new construction in response to a demand shock, suggesting that this index reflects meaningful variation in elasticity of housing supply across locations.

2.1. Measuring housing supply regulations across metropolitan areas

Government regulations can influence residential construction in numerous ways. Land use policy is generally controlled by local governments, and the political

[2] Although a number of studies have examined the correlation between house prices and migration, only a few specifically address the effect of the housing supply on local labor markets. One example is Case (1991), who discusses this issue in the context of rising labor demand in Boston in the 1980s. Also, Bover et al. (1989) analyze the effect of regional housing market constraints on migration flows in the UK. Neither of these papers attempts to identify the effect of housing supply constraints separately from housing demand, as I will do in this paper.

[3] In this regression and in all of the analysis that follows, metropolitan areas are defined using the 1999 Census definitions of PMSAs and NECMAs. County-level data are aggregated to the metropolitan level so that metropolitan area boundaries can be kept fixed over time. See Appendix A for details.

[4] Some mechanisms that allow labor markets to adjust when the housing supply remains fixed include changes in the unemployment rate, labor force participation and the housing vacancy rate. However, prior studies show that these margins of adjustment appear to be small. Blanchard and Katz (1992) find that labor force participation and unemployment account for about 50 percent of the impact of a labor demand shock in the first year, and that these factors become less important over longer time horizons. Hwang and Quigley (2006) find that vacancy rates are only weakly related to labor market conditions in metropolitan areas. Glaeser et al. (2006) also show that changes in vacancy rates and household size only explain a small fraction of the variance of changes in population across cities.

[5] Although one could theoretically identify this elasticity from a regression of house prices on quantity and instrumenting with shocks to housing demand, in practice it is difficult to find exogenous instruments that are strong enough to yield precise estimates. Along these lines, in Saks (2004) I use an instrumental variable strategy to estimate the responses of construction and house prices to a labor demand shock in 131 individual metropolitan areas. I find substantial variation in the elasticity of housing supply across locations. Other related research includes Green et al. (1999), who estimate MSA-specific elasticities from a reduced form regression of prices on quantities, and Evenson (2002), who estimates local elasticities from changes in employment. However, neither of these studies identifies the slope of the supply curve from exogenous changes in housing demand.

Table 1
Survey measures of housing supply regulation

Source	Time period	# of survey questions	Type of jurisdiction	# of MSAs
Wharton Urban Decentralization Project	1988–1990	8	Metropolitan areas	60
International City Management Association	1984	4	Cities	155
Fiscal Austerity and Urban Innovation Project	1983–1984	1	Cities	62
Regional Council of Governments	1975–1978	1	Metropolitan areas	48
National Register of Historic Places	Prior to 1980	2	Counties	318
American Institute of Planners	1976	8	States	318

environment in each individual municipality shapes the degree of regulation and the form that it takes (Fischel, 1985). Thus, these laws vary widely across local jurisdictions, making it difficult to summarize the extent of regulation in a manner that is comparable across a large number of locations. Consequently, the types of regulation and number of locations covered by empirical surveys of land use policies are generally limited. In order to develop a more comprehensive view of the extent of regulation in a large number of locations, I combine information from six sources into a single index. Table 1 lists each of the sources that are used, the time period and type of geographic areas that are covered. This section contains a brief summary of each of the index's components and how it is constructed, and full details can be found in the Appendix.

The first source is a survey conducted by the Wharton Urban Decentralization Project (Wharton) in the late 1980s, which surveyed city planners and local county officials concerning local development regulations. I aggregate the answers to nine individual survey questions including the fraction of zoning permits approved, the length of time for permits to be processed, and the importance of growth management policies.[6] Due to the relatively large number of jurisdictions covered and questions asked, this survey has been the primary source of information cross-metropolitan area differences in housing supply regulation (Malpezzi, 1996; Mayer and Somerville, 2000).

The second survey was undertaken by the Fiscal Austerity and Urban Innovation project, which asked city managers to rate the importance of restricting construction in order to limit population growth. Thus, this survey reflects the severity of growth controls, which have become an increasingly popular tool for managing development during the past two decades (Dubin et al., 1992; Seidel, 1978). A third constraint on residential construction stems from the desire to protect certain areas from development for historic reasons. For example, many cities have a historic district to preserve buildings of historic significance or the historic character of the neighborhood. Even in suburban or rural areas, development can be limited by the existence of historic landmarks, archeological sites, and old battlegrounds. To create a measure of the extent that residential construction might be limited by concerns about historic preservation, I collect information on the fraction of land area reserved for historic districts or sites and the number of historic buildings and structures per square mile from the National Register of Historic Places (NRHP).

Although the majority of land use regulations are imposed by local jurisdictions, state governments also play a role in influencing patterns of development. For example, a state may restrict development in certain areas, mandate certain types of land use planning or environmental impact analyses, or control the development of new towns. Therefore, the fourth type of supply constraint I consider is based on a study of state regulations conducted by the American Institute of Planners (AIP).

Finally, I include responses from two surveys of local government officials that concern the general degree of housing supply regulation in each location. The first is a survey of the Regional Council of Governments, which asked officials to estimate the fraction of suburban land area made unavailable for development as a result of government regulation. The second is a survey conducted by the International City Management Association, which asked members how frequently zoning or environmental regulations were altered to facilitate economic development.

Among these six sources of information, only the AIP survey of state regulations and the index of historic preservation can be calculated for every metropolitan area in the United States. The four other surveys were

[6] It is possible that the permit-approval process would take longer in places experiencing a large unanticipated increase in housing demand. Furthermore, if strict regulations discourage potential builders from applying for building permits, the fraction of permits approved would not necessarily be lower in locations with a high degree of regulation. This would tend to dampen the variation of permit approvals across locations, making it more difficult to observe any effect of these regulations on housing markets.

Table 2
Ten most and least regulated metropolitan areas

Most regulated		Least regulated	
Index value	Name	Index value	Name
2.21	New York, NY	−2.40	Bloomington–Normal, IL
2.10	San Francisco, CA	−1.96	Buffalo–Niagara Falls, NY
1.89	Sacramento, CA	−1.65	Nashville, TN
1.84	Charleston–North Charleston, SC	−1.50	Owensboro, KY
1.73	Riverside–San Bernardino, CA	−1.48	Joplin, MO
1.65	San Jose, CA	−1.37	Pueblo, CO
1.60	San Diego, CA	−1.36	Champaign–Urbana, IL
1.51	Santa Barbara–Santa Maria–Lompoc, CA	−1.32	Oklahoma City, OK
1.48	Seattle–Bellevue–Everett, WA	−1.26	Dayton–Springfield, OH
1.23	Gary, IN	−1.23	Richmond–Petersburg, VA

Notes. The index of housing supply regulation is calculated by combining information from each of the sources listed in Table 1. The index is scaled to have a mean of 0 and a standard deviation of 1, and is increasing in the degree of regulation. See Appendix A for details.

conducted in a limited number of metropolitan areas, and the geographic areas covered by each one are considerably different—only 17 metropolitan areas have information from all six sources. Rather than limiting the analysis to these places, I use a set of OLS regressions to predict missing values of each survey from observed values of the other surveys.[7] The final index includes all of the locations with non-imputed information for at least four of the six sources, for a total of 83 metropolitan areas.

Table 2 lists the most and least regulated metropolitan areas according to this measure, and Table A.2 in the appendix reports the index values for all locations. It comes as no surprise to find that San Francisco, CA, Seattle, WA and New York, NY are among the most highly regulated locations. Areas with the least amount of regulation include Nashville, TN, Pueblo, CO and Champaign, IL. The two panels of Fig. 1 show that places with a large amount of regulation have experienced higher housing price inflation and lower residential construction from 1980 to 2002.[8]

2.2. *Effects of regulation on the housing market*

Although the correlations in Fig. 1 are suggestive, metropolitan areas differ along many unobservable dimensions that can make it difficult to sort out the effect of housing supply from other factors. Instead of using this cross-sectional comparison, the impact of the housing supply can also be found by examining relative housing market dynamics *within* metropolitan areas over time. In particular, suppose that the elasticity of housing supply is a function of regulation (r_i) and other factors (u_i):

$$p_{it} = \theta_i h_{it} + \eta_{it}, \tag{1}$$

$$\theta_i = \theta_0 + \pi r_i + u_i \tag{2}$$

where p_{it} is the logarithm of the house price and h_{it} is the logarithm of the housing stock in metropolitan area i at time t. The parameter θ_i represents the inverse of the elasticity of housing supply, as a larger value of θ_i indicates that a given increase in the housing stock would lead to a larger rise in house prices. Through the parameter π, differences in housing supply regulation generate variation in the elasticity of housing supply across locations. Substituting Eq. (2) into (1) shows that π can be estimated from the interaction of the housing stock with these regulations:

$$p_{it} = \theta_0 h_{it} + \pi r_i h_{it} + u_i h_{it} + \eta_{it}. \tag{3}$$

Thus, a positive value of π indicates that locations with stricter regulation have a more inelastic housing supply. Because changes in the housing stock are likely to be correlated with the error term in this equation, I estimate a fixed-effects version of Eq. (3):

$$p_{it} = \theta_0 h_{it} + \pi r_i h_{it} + x_i + d_t + \varepsilon_{it}. \tag{4}$$

By including year and metropolitan area fixed effects in the regression, the effect of regulation is identified from deviations of construction and housing price growth from location-specific and time-specific averages. Thus, the estimates will not be biased by omitted factors like geographic amenities that are constant over time.

[7] In particular, I impute the value of a given source S in location X using the fitted values from an OLS regression of S on all of the other sources that are available for location X.

[8] This figure and the regression in Section 2.3 omit low-demand metropolitan areas because the supply response to an increase in housing demand will be different in locations with sustained low levels of housing demand (Glaeser and Gyourko, 2005). See Appendix A for the method used to identify these locations.

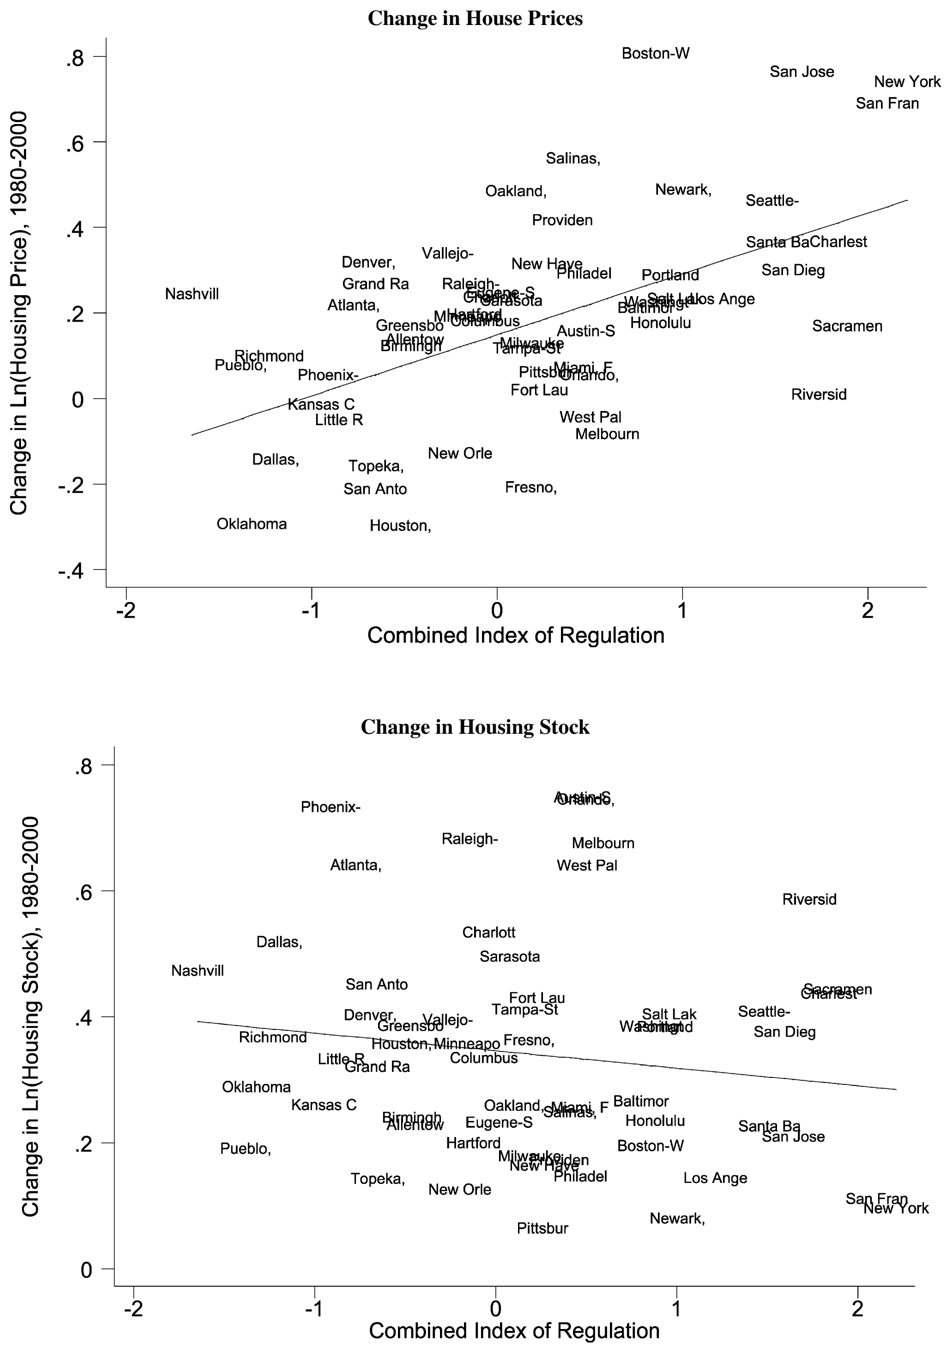

Fig. 1. Combined Index of Housing Supply Regulation and Growth in Growth in Housing Prices and Quantities, 1980–2000.

Even with these fixed effects, changes in the housing stock are likely to be correlated with the error term ε_{it}. Therefore, I replace h_{it} with a variable reflecting shocks to housing demand. To measure these shocks, I use annual changes in labor demand as predicted from the industrial composition of each metropolitan area (Bartik, 1991). The idea behind this strategy is that firms in the same industry face similar conditions in the product market, and thus are likely to have similar demands for workers, irrespective of geographic location. Assuming that every firm would like to hire workers at a rate equal to the change in employment in its industry, employment growth for each metropolitan area can be predicted as a weighted average of national indus-

Table 3
Effects of demand shocks and housing supply regulation on metropolitan area housing markets

Dependent variable	Percent change in housing stock	Percent change in house prices
Demand shock	0.254**	0.807
	(0.096)	(0.458)
Demand * regulation	−0.027*	0.120
	(0.016)	(0.099)
Year fixed effects	yes	yes
MSA fixed effects	yes	yes
# MSAs	58	58
# Observations	1276	1276

Notes. Each column shows the coefficients from a separate OLS regression where the dependent variable is changes in the housing stock (left panel) or changes in housing prices (right panel). Regulation is the index of housing supply regulation described in the text. Standard errors are reported in parentheses and are clustered by state.
* Significance at the 10% level.
** Idem, 5%.

try growth rates, where the weights are determined by the industrial composition of the area.[9] For example, predicted labor demand in areas with a large share of automobile manufacturing plants will be high when the automobile industry is hiring more workers nationally relative to firms in other industries.

Table 3 reports the results of estimating a regression of annual changes in the logarithm of house prices and the housing stock on the labor demand shocks and an interaction with the index of housing supply regulation.[10] On average, a one percent increase in labor demand is associated with a 0.25 percent increase in the housing stock and a 0.8 percent increase in housing prices. Consistent with the notion that regulations decrease the elasticity of housing supply, the interaction of the labor demand shocks with the index of housing supply regulation is negative for housing quantities and positive for housing price inflation. Thus, this index of regulation appears to reflect meaningful differences in the elasticity of housing supply across locations.[11]

3. A model of housing and labor markets

Through their effect on construction and house prices, constraints on the supply of housing will also impact local labor markets. In this section, I develop a simple framework to illustrate this connection. I begin with the basic model of regional labor markets from Blanchard and Katz (1992) and extend it to incorporate local housing markets. The economy is made up of a large number of metropolitan areas, each indexed by i. In each place, the marginal product of labor declines with the level of employment, so that the demand for labor is downward sloping:

$$w_{it} = -\delta n_{it} + z_{it} \tag{5a}$$

where w_{it} is the wage and n_{it} is employment in area i at time t. All variables are measured in logs and reflect deviations from the national average. The variable z_{it} represents shifts in the labor demand curve, and is assumed to contain both a unit root and drift component:

$$z_{it} - z_{it-1} = x_i^d + \varepsilon_{it}^d. \tag{5b}$$

The fixed city-specific factors in this equation capture any local characteristics that cause labor demand to differ systematically across locations. For example, each metropolitan area produces a different combination of goods, and so we would expect the labor demand curves in each location to shift differentially as the relative demand for goods changes over time. These factors imply that cities are likely to grow systematically at different rates, a prediction that fits the patterns of local employment growth well during the post-WWII period (Blanchard and Katz, 1992). A shock to the relative productivity of an area is reflected in the idiosyncratic term ε_{it}^d. Because the position of the demand curve follows a random walk, the effects of all labor demand shocks are permanent.[12]

9 I exclude employment in the construction industry because it is likely to be correlated with the elasticity of housing supply. To ensure that these employment growth predictions are not related to local conditions, the calculation for each metropolitan area subtracts industry employment in that individual area from the national industry total. In other words, the shocks are based on industrial employment growth outside of the specific metropolitan area in question. I also subtract national employment growth from these adjusted industry growth rates so that an increase in aggregate demand will not lead to a change in predicted labor demand. Thus, these shocks reflect purely relative changes in the local demand for labor. See Appendix A for the exact formula used for calculating these shocks.

10 Annual changes in the housing stock are calculated from information on the number of residential building permits issued in each county and the size of the housing stock in the 1980, 1990 and 2000 Census. See Appendix A for details.

11 These results are robust to including interactions of the labor demand shock with the following measures of differences in productivity and the supply of land across locations: metropolitan area age, the logarithm of January temperature, density of housing units in 1980, the fraction of total area taken up by water, and an indicator for negative demand shocks.

12 The unit root in shocks to labor demand captures the notion that both changes in relative product demand and technological progress are highly persistent. Because variation in local demand shocks is largely related to differences in the industrial composition of each location, local demand shocks are also likely to be highly persistent.

The supply of labor in each area is determined by the size of the population. There is no adjustment through changes in hours or unemployment, meaning that the supply of labor in the short run is completely inelastic. Over time, workers respond to relative differences in wages and housing prices by moving between locations. Migration into an area increases with the relative level of wages (w_{it}) and decreases with the relative level of housing prices (p_{it}). Migration also depends on fixed location-specific factors (x_i^s), such as weather, that cause some areas to be permanently more attractive than others.

$$n_{it} - n_{it-1} = \beta w_{it-1} - \gamma p_{it-1} + x_i^s + \varepsilon_{it}^s. \tag{6}$$

Because house prices affect migration, the state of the housing market will impact equilibrium wages and employment. To model the housing market, I assume that everyone in the population works and that all workers must live in a separate house, so that housing demand, and therefore the equilibrium size of the housing stock, is equal to employment. On the supply side of the housing market, I use an expression similar to Eq. (1) where the size of the housing stock is equal to the size of the population:

$$p_{it} = \theta_i n_{it} + x_i^p + \varepsilon_{it}^p. \tag{7}$$

As discussed in Section 2, the parameter θ_i reflects the inverse of the elasticity of housing supply. A high value of θ_i means that the housing supply is more inelastic, as a given increase the size of the housing stock is translated into higher prices. The supply of housing is also allowed to depend on any fixed city-specific factors (x_i^p) factors that might create persistent differentials in average housing prices across locations.

To see the effect of an increase in labor demand on each of the variables in the model, Eqs. (5)–(7) can be rewritten to express each variable as a function of its own previous values and the shocks:

$$\begin{aligned} n_{it} - n_{it-1} = {} & (1 - \beta\delta - \gamma\theta_i)(n_{it-1} - n_{it-2}) \\ & + \beta x_i^d + \beta\varepsilon_{it-1}^d + \varepsilon_{it}^s - \varepsilon_{it-1}^s + \gamma\varepsilon_{it}^p \\ & - \gamma\varepsilon_{it-1}^p, \end{aligned} \tag{8a}$$

$$\begin{aligned} w_{it} = {} & (1 - \beta\delta - \gamma\theta_i)w_{it-1} + x_i^d - \delta x_i^s + \gamma\delta x_i^p + \varepsilon_{it}^d \\ & + \gamma\theta_i\varepsilon_{it-1}^d - \delta\varepsilon_{it}^s + \gamma\delta\varepsilon_{it-1}^p + \gamma\theta_i z_{it-2}, \end{aligned} \tag{8b}$$

$$\begin{aligned} p_{it} = {} & (1 - \beta\delta - \gamma\theta_i)p_{it-1} + \beta\theta_i x_i^d + \theta_i x_i^s - \beta\delta x_i^p \\ & + \beta\theta_i\varepsilon_{it-1}^d + \theta_i\varepsilon_{it}^s + \varepsilon_{it}^p - (1 + \beta\delta)\varepsilon_{it-1}^p \\ & + \beta\theta_i z_{it-2}. \end{aligned} \tag{8c}$$

The initial impact of an increase in labor demand (ε_{it}^d) is an increase in wages. There is no immediate change in employment or house prices because migration only depends on lagged values of house prices and wages, making the labor supply infinitely inelastic in the first period.[13] In the second period, higher relative wages cause workers to migrate into the location that experienced the demand shock. This increase in population creates additional demand for housing, and so house prices rise as well. The response of house prices depends on the elasticity of housing supply, as more inelastic areas (higher θ_i) experience a larger increase in housing prices. Because the initial shock has a permanent effect on labor demand, the shock continues create an additional upward pressure on wages and house prices. These lagged responses of wages and prices also depend on the elasticity of housing supply, because migration (and therefore the labor supply adjustment) depends on the amount of new construction. After the first period, employment rises as migrants are attracted by the increase in wages. In the second period, the ratio of the increase in house prices to the increase in employment in response to a 1-unit increase in the labor demand shock (ε_{it}^d) is $\beta\theta_i/\beta = \theta_i$. In the empirical analysis below, I will use this relationship to infer the value of θ_i for each metropolitan area.

The model shows that the elasticity of housing supply impacts the labor market through house prices and the resulting migration response. Higher values of θ_i create more persistence in all three variables, so that the effect of any shock takes longer to dissipate. Even though the initial shock leads to a permanent increase in labor demand (z_{it}), out-migration in response to higher house prices partly offsets the upward pressure on wages and employment. In the long run, a new equilibrium is reached where the ratio of wages to house prices equals γ/β in all locations. Because this ratio is the same in all locations, there is no further incentive to move, and consequently no further growth in employment. Thus, in the long run the *change* in employment converges to zero even as the labor demand curve continues to shift out. At this new equilibrium, the final *level* of employment in the location that experienced the demand shock is higher than its initial level. However, the initial boost to employment is diminished

[13] In reality, the migration decision is likely to be a function of expected future conditions and not just on past experience. However, if expectations are backward-looking, then estimates of the model parameters will be a function of the true parameters and the discount rate of households (Gallin, 2004). Assuming that households in different locations have a similar discount rate, estimates of the relative differences in the elasticity of housing supply will not be affected by the simplifying assumptions made in this model.

by out-migration to locations with lower house prices. Assuming that all of the fixed effects are zero (which corresponds to a metropolitan area with average values of wages, house prices, and employment growth) and that there are no shocks to housing or labor supply, the long run levels of employment, wages and housing prices after a 1-unit increase in ε_{it}^{d} are[14]:

$$\hat{n}_i = \frac{1}{\theta_i}\hat{p}_i = \frac{\beta}{\gamma\theta_i + \beta\delta},$$

$$\hat{w}_i = \lim_{t\to\infty}(1-\gamma\theta_i-\beta\delta)^t + \gamma\theta_i\sum_{j=0}^{j=t}(1-\gamma\theta_i-\beta\delta)^j \to 0 + \frac{\gamma\theta_i}{\gamma\theta_i+\beta\delta},$$

$$\hat{p}_i = \lim_{t\to\infty}\beta\theta_i\sum_{j=0}^{j=t}(1-\gamma\theta_i-\beta\delta)^j \to \frac{\beta\theta_i}{\gamma\theta_i+\beta\delta}. \quad (9)$$

These expressions reveal that the elasticity of housing supply alters the impact of a labor demand shock on the long-run evolution of all of the endogenous variables in the model. Areas with a less responsive housing supply will end up with higher wages, higher house prices, and a lower level of employment than a location with a more elastic housing supply.

One counter-intuitive prediction of this model is that the labor demand shock will have a lasting impact on wages, as long as house prices are also permanently higher. By contrast, prior studies have shown that wage differentials across locations tend to converge over time (Barro and Sala-i-Martin, 1991; Blanchard and Katz, 1992). In the original model developed by Blanchard and Katz, relative wages converge because firms relocate to areas with lower wages. The model presented here rules out this response because relative wage convergence would mean that relative house prices must converge as well. Because the position of the housing supply curve is fixed, the only way that house prices could return to their initial level is if the housing stock (and therefore the level of employment) returned to its initial level as well. Thus, relative wage convergence in this model would imply that an increase in local labor demand would be completely offset in the long run as firms evade higher wages.

To address this issue, firm mobility could be included in the model as long as the housing supply curve is also allowed to shift in response to relative house price differentials. In particular, if higher house prices encourage entry into the construction industry, the housing supply curve would shift out in response to an initial increase in house prices. In this case, an increase in labor demand would have no long-run effect on either relative wages or house prices, because relative wage and price differentials would be offset by firm migration in both the construction industry and in other sectors. Instead, the effect of the demand shock would result only in a higher level of employment and a larger housing stock. Although this model has more realistic long-run predictions for wages, I focus on the model without firm entry and exit in either the product or housing markets for simplicity. However, it should be kept in mind that empirically, the wage and price differentials predicted by this model are not likely to persist in the very long run.

4. Estimating the effect of the housing supply on local labor markets

4.1. Baseline short- and long-run effects

The model discussed in the previous section illustrates how the effect of an increase in labor demand on employment, wages and house prices depends on the elasticity of housing supply. To assess these predictions empirically, I trace out the effect of an increase in labor demand on metropolitan area housing and labor markets. As described in Section 2, I follow Bartik (1991) and calculate shocks to labor demand arising from differences in the industrial composition of metropolitan areas interacted with national shocks to industrial employment growth. To observe the effect of these shocks on the dynamics of employment, wages

[14] Because the initial shock ε_{it}^{d} returns to zero after the initial period, changes in employment converge to zero: $\lim_{t\to\infty}\beta(1-\gamma\theta_i-\beta\delta)^t \to 0$. However, because the shock ε_{it}^{d} raises z_{it} permanently, the long-run levels of wages, prices and employment remain above their initial levels. The long-run levels of wages and prices are determined by Eqs. (8b) and (8c), and the level of employment can be found by plugging the long-run level of prices into the housing supply equation (7). These solutions require the assumption that $|1-\beta\delta-\gamma\theta_i| < 1$. In a survey of estimates of the elasticity of labor demand, Hamermesh (1993) concludes that the aggregate elasticity of labor demand is in the range of -0.15 to -0.75, which implies that δ, the elasticity of wages with respect to labor supply, is between -0.12 and -0.6 (see pp. 26–29). DiPasquale (1999) reviews the literature on the elasticity of housing supply and concludes that the aggregate housing supply is elastic: most estimates range between 1.5 and 5. Therefore θ, which is the inverse of the elasticity of housing supply, is likely to be less than 1 for the typical metropolitan area. The magnitudes of β and γ are more difficult to determine because they are related the parameters of an individual's utility function. Allowing the fixed effects to differ from zero would not qualitatively change these solutions. Rather, Eq. (9) would express the long-run deviation of each variable from its hypothetical evolution had there been no labor demand shock.

and house prices, I estimate the following 3-variable Vector Auto-Regression (VAR):

$$Y_{it} = \begin{bmatrix} \Delta n_{it} \\ w_{it} \\ p_{it} \end{bmatrix} = B_1 Y_{it-1} + B_2 Y_{it-2} + B_1^r Y_{it-1} reg_i + B_2^r Y_{it-2} reg_i + C\hat{\varepsilon}_{it}^d + C^r \hat{\varepsilon}_{it}^d reg_i + D_i + D_t + V_{it}. \tag{10}$$

This system of equations expresses the change in the logarithm of employment, the logarithm of wages, and the logarithm of house prices each as a function of two of its own lags, two lags of the other endogenous variables, and the contemporaneous labor demand shocks ($\hat{\varepsilon}_{it}^d$).[15] This system of equations reflects a reduced-form version of Eq. (8), where the lagged endogenous variables other than the lagged dependent variable reflect unobservable shocks to the housing supply or labor supply (ε^{p} and ε^{s}). This autoregressive VAR can also be written in its moving-average representation, in which each variable is expressed as a function of only contemporaneous and lagged values of the shocks.[16] The impulse response functions found by estimating the autoregressive form of the VAR reflect the moving-average representation of the model. Thus, this VAR describes the dynamics of employment, wages, and house prices in response to the underlying shocks.

As suggested by Eq. (2), I capture geographic variation in the coefficients on the demand shock and the endogenous variables (which results from geographic differences in θ_{i}) by interacting these variables with the index of regulation. Each equation includes year and metropolitan area fixed effects, and is estimated using annual data from 1980 to 2002. Thus, the results are identified from the behavior of relative wages, employment and housing prices within each metropolitan area over time. This method is more useful than estimating the effects from a single cross-section because locations differ along many unobservable dimensions that can easily confound a cross-sectional comparison of locations. The final sample is a balanced panel comprising a total of 72 metropolitan areas, which includes all of the metropolitan areas with a value of the housing supply regulation index and complete information on house prices from 1980 to 2002.

Table 4
Initial effects of a one-percent increase in labor demand on employment, wages and house prices

	Demand shock	Demand shock × Regulation
Employment growth	1.04**	−0.06**
	(0.20)	(0.02)
Ln(Wage)	0.68**	0.07**
	(0.11)	(0.03)
Ln(House price)	0.13	0.15
	(0.22)	(0.10)
Implied $\theta = \frac{1}{\text{elasticity of housing supply}}$	0.13	0.29
	(0.20)	(0.22)
Chi^2 test that θ is equal across areas	2.14	
	(0.14)	

Notes. Cells in the top three rows show the coefficients on the labor demand shock from the VAR described in the text, with standard errors in parentheses (clustered by state). Regressions are estimated using annual data from 1980 to 2002 and include time and MSA fixed effects. In the first column, θ is calculated as the ratio of the third row to the first row. The second column shows the ratio of the sum of the base effect and interaction term in the third row to the same sum in the first row. The final row shows the Chi^2 statistic testing that the two estimates of θ are equal, with p-value in parentheses.
* Significance at the 10% level.
** Idem, 5%.

The initial effects of the labor demand shock on each of the three endogenous variables in the system are shown in Table 4. The coefficient on the demand shock in the employment growth equation is 1.04, which implies that in the average metropolitan area, there is a strong relationship between increases in labor demand and increases in employment. This result is reassuring because it indicates that the method of calculating labor demand shocks yields accurate predictions of actual increases in labor demand. In the average metropolitan area, an increase in labor demand is also associated with higher wages, but only a small and insignificant change in house prices.

The interaction terms in the second column show how housing supply regulations alter the responses of employment, wages and house prices. In areas where the housing supply is more constrained, an increase in labor demand leads to higher house prices. This effect is accompanied by a higher level of wages and a smaller increase in employment, which is a sign that migration into these areas is constrained. The fourth row of the table shows implied estimates of θ, the inverse of the elasticity of housing supply. These estimates are obtained by taking the ratio of the response of house prices to a change in labor demand ($\beta\theta_{\mathrm{i}}$) to the re-

[15] I allow for only two lags of each variable because the time dimension of the panel is relatively short, extending for a total of 22 years. Adding a third lag does not change the results substantially, as the coefficients on the 3rd lags are generally small and insignificantly different from zero.

[16] The moving-average representation of the model can be found by substituting the lagged dependent variables recursively into Eq. (8).

sponse of employment growth to the same labor demand shock (β). The results are consistent with a lower elasticity in highly constrained areas, as a one standard deviation increase in the housing supply regulation index is associated with an estimate of θ that is twice as high. However, the standard errors are large enough that the two estimates are not significantly different from one another.[17]

To examine the long-run impact of housing supply regulations, Fig. 2 shows the impulse response functions from the VAR. The solid lines show the response of each endogenous variable to a one percent increase in labor demand in an area at the 25th percentile of the housing supply regulation index (equivalent to the index value of Denver, CO). I calculate these long-run effects by setting the response of each variable in the first period equal to the estimated coefficient shown in the first column of Table 4 plus the interaction term multiplied by the value of the index at the 25th percentile (-0.69). In the second period the labor demand shock is set equal to zero, but the endogenous variables continue to evolve through the lagged endogenous variables and the interactions of these lagged variables with the degree of housing supply regulation. I convert predicted employment growth to a level by assuming the logarithm of the initial level of employment is zero.

Initially, the shock leads to increases in employment and wages, with only a small change in house prices. The level of employment rises for the next few years and then declines, converging to a long-run effect of about 1 percent after 15 years. Relative wages decline more gradually, taking about 25 years to return to their initial level. In the first few years following the employment shock, house prices rise by about 1.8 percent. This increase is due to a positive effect of lagged employment growth on house prices. Since the housing market is not perfectly elastic, the inflow of migration that was generated by the initial the demand shock creates excess housing demand, causing house prices to rise. Over time, house prices fall below their initial level before converging to the long-run equilibrium. This result illustrates that the housing supply tends to react to changes in local market conditions with a lag, leading to oscillation of prices as the housing supply first undershoots, and then exceeds, the long-run level of housing demand.

The dashed lines in the figure show the impulse response functions for an area at the 75th percentile of the housing supply regulation index (equivalent to Newark, NJ). As predicted by the model, the demand shock has a smaller impact on employment and a larger impact on wages and house prices. After the large, initial boost to increase in house prices, prices continue to rise as labor migration raises the demand for housing. Wages are also persistently higher for many years, before eventually converging back to their initial level. By contrast, the level of employment never reaches the corresponding level of employment in the less constrained location. The long-run impact of a one percent increase in labor demand results in only a 0.9 percent increase in employment, instead of the one percent increase found in less constrained areas. However, using a bootstrap technique to calculate confidence intervals around these estimates, 90 percent of the simulated estimates fall in the range of [0.83, 1.23] for an elastic location and [0.68, 1.10] for the inelastic location. Thus, these long-run employment effects are not significantly different from one another.

4.2. *Robustness checks*

The effects shown in Fig. 2 rely on the assumption that the labor demand shocks are uncorrelated with the error term in each of the three estimating equations. However, these estimates will be biased if the effect of a demand shock is smaller in places with more regulation for reasons other than housing supply constraints. For example, if places with a high degree of housing market regulation also have a stronger regulatory environment more generally, the negative impact of housing supply regulations on employment will be overstated. On the other hand, factors like weather might cause a labor demand shock to lead to a larger employment response in places where it is more pleasant to live. If this geographic amenity is correlated with a stronger degree of housing supply regulation, the estimated employment effects will be biased downward. A third source of concern is the age of the metropolitan area. If firms in older cities have a less productive capital stock, they will hire fewer workers than firms in the same industry that are located in newer cities. In this case, the measured demand shocks will overstate the actual amount of labor demand in older cities. If there are also more housing supply regulations in older cities, then the estimated effect of regulation will be biased.

To address these issues, I re-estimate the VAR including a complete set of interactions of all of the right-hand side variables with observable measures of the factors discussed above. The four variables that I consider are an index of the regulatory environment at the state level, the extent of unionization in the metropolitan

[17] These estimates are not far different from DiPasquale's (1999) survey of the literature on the housing supply, who concluded that the aggregate elasticity of housing supply ($1/\theta_i$) is between 1.5 and 5.

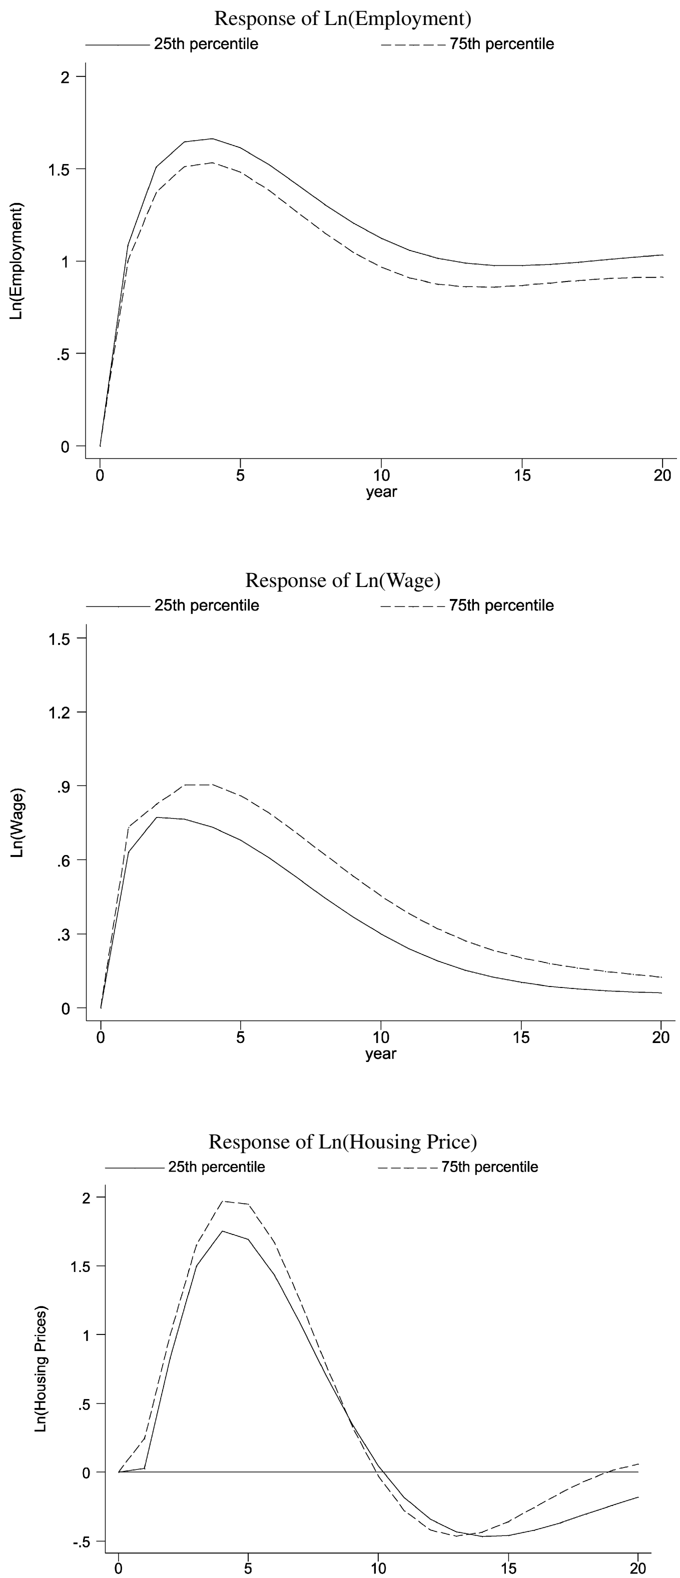

Fig. 2. Long-run response to a 1-percent increase in labor demand evaluated at the 25th and 75th percentiles of the housing supply regulation index.

Table 5
Long-run employment responses to a one-percent increase in labor demand controlling for other factors

Control variable	Housing supply regulation		Control variable (named in row)	
	25th percentile	75th percentile	25th percentile	75th percentile
January temperature	0.90	0.79	0.88	0.80
MSA age	0.87	0.71	0.81	0.79
Unionization	1.02	0.88	0.92	0.97
State regulation	1.00	0.93	1.03	0.92
Fraction water	1.07	0.97	1.05	1.02
Housing density	0.98	0.88	0.94	0.84

Notes. Each cell shows the level of employment after 20 years in response to a 1 percent increase in labor demand. Each row is from a separate VAR as described in the text. In each VAR, both the index of housing supply regulation and the control variable named in the row are fully interacted with the labor demand shock and all lagged endogenous variables. Impulse response functions are calculated at the 25th and 75th percentiles of housing supply regulation, and then at the 25th and 75th percentiles of the control variable named in each row. Variable definitions can be found in Appendix A.

area (also as a measure of regulation in the labor market), the logarithm of January temperature, and the age of the metropolitan area.[18] I introduce each of these four variables in a separate VAR because including them simultaneously would require estimating too many new parameters. The interactions of these variables with the labor demand shock and the lagged endogenous variables allow the labor and housing market dynamics in each location to depend on each new control variable, as well as on the degree of housing supply regulation. As before, I estimate the level of employment after 20 years in response to a one percent increase in labor demand for a metropolitan area at the 25th and 75th percentile of the distribution of housing supply regulation (columns 1 and 2 of Table 5). Each control variable is normalized to have a mean of zero, so these results can be interpreted as occurring at the mean of the other control variables. Similarly, columns 3 and 4 of Table 5 report the long-run effect on employment in metropolitan areas at the 25th and 75th percentiles of each other control variable.

In each case, the estimated impact of housing supply regulation is essentially the same as in the baseline specification, as employment growth is about 10 percentage points lower in areas that are relatively more constrained. The majority of the estimated effects of the other control variables also go in the expected direction. For example, an area at the 75th percentile of the state regulatory index has an 11 percentage point $(0.92 - 1.03 = -0.11)$ smaller employment response than an area at the 25th percentile.[19]

Another source of concern is that the index of housing supply regulation might be correlated with geographic factors that make the supply of housing more inelastic in some locations relative to others. Therefore, the final two rows of Table 5 show the results of specifications that include interactions with two variables that proxy for the supply of land: the share of total area taken up by water, and the number of housing units per square kilometer in 1980. Although these variables also appear to have a negative impact on employment, the estimated effect of regulation is unchanged when they are included in the regression.

A final issue is that local governments may impose stricter regulations in locations where the expected future demand for housing is higher. These expectations could arise from location-specific trends in either productivity or consumption amenities. Although the metropolitan area fixed effects account for differences in average demand and amenities across locations, they do not allow for the possibility of area-specific trends in these factors. Expected future demand conditions should be positively correlated with both house prices and employment, biasing the positive interaction between regulation and house prices upward and the negative interaction between employment and regulation toward zero. Thus, although these results may overstate the effect of regulation on house prices, they may understate negative impact on employment. However, the results are similar if MSA-specific time trends—which should at least partly capture these unobservable trends—are included in the regression, suggesting that the magnitudes of these biases are not large.

As a final illustration of the variation in the elasticity of housing supply across locations, I use the VAR results from Table 4 to calculate impulse response functions at each of the observed values of the regulation index. Fig. 3 shows the distribution across locations of the predicted level of employment after 20 years. The response to a one percent increase in labor demand in many of the metropolitan areas is close to one percent, which emphasizes that the housing supply in many locations is fairly elastic in the long run. However, the distribution of these employment predictions is highly skewed, with a significant number of metropolitan areas

[18] See Appendix A for definitions and sources of these variables.

[19] Surprisingly, warmer temperatures and less unionization appear to lead to lower employment in the long run. However, the short-run effects of these variables are more sensible, as employment is substantially higher in warmer and less unionized places in the first 11 years after the shock.

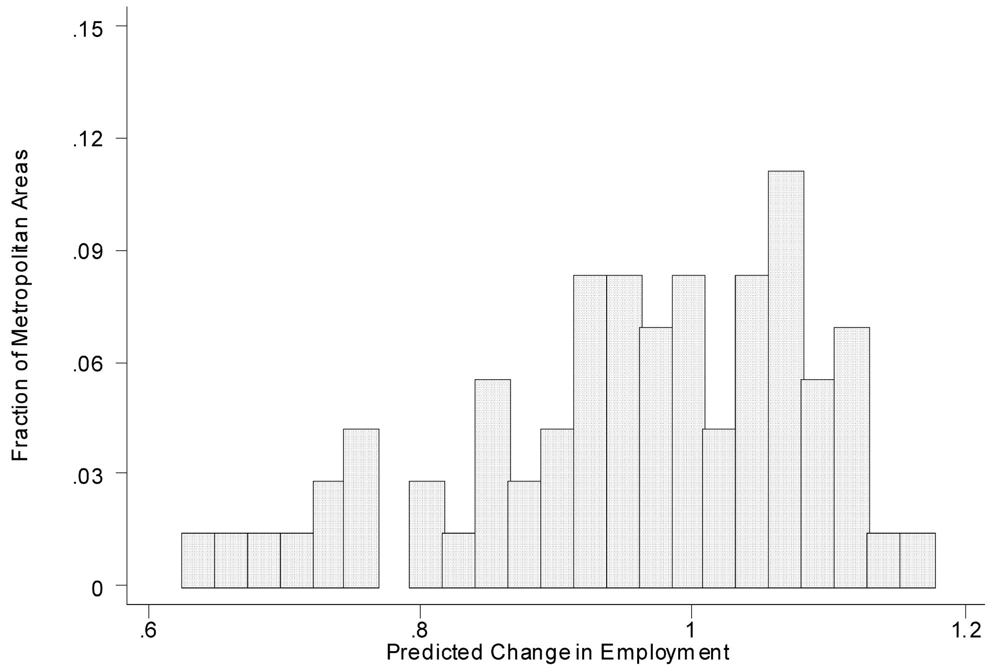

Fig. 3. Distribution of long-run employment responses to a 1-percent increase in labor demand.

in the left-hand tail. For example in New York, which was estimated to have the highest degree of regulation, a one percent increase in labor demand leads to only a 0.65 percent increase in employment. The 90 percent confidence band around this estimate encompasses the interval [0.42, 0.96], which suggests that this estimate is significantly less than one. The long-run employment response is significantly less than one for the 12 most constrained locations, or about 16 percent of the metropolitan areas in the sample. Thus, although housing supply constraints may not have had a large impact on employment in the majority of metropolitan areas, these results suggest that they have been a limiting factor in the locations where regulation has been the most severe.

5. Conclusion

Land use restrictions and other government regulations have a substantial impact on housing and labor market dynamics in metropolitan areas across the United States. These regulations lower the elasticity of housing supply, consequently altering relative differences in house prices and patterns of labor migration across locations. As a result, employment growth is lower in places where the housing supply is more constrained.

The empirical analysis in this paper showed that locations with a larger degree of housing supply regulation experience less residential construction and larger increases in house prices in response to an increase in labor demand. Moreover, housing supply regulations have a lasting effect on metropolitan area employment. In the long run, an increase in labor demand results in considerably lower employment in metropolitan areas with a low elasticity of housing supply. These results demonstrate that the interaction between the housing supply and local labor markets is an important determinant of regional patterns of employment growth.

The impact of housing supply regulations likely extends beyond the effects on employment, wages and house prices discussed here. For example, as places with a large degree of regulation experience rising house prices, the composition of the population within these metropolitan areas may change. Because young people and minorities have a higher propensity to move (US Census Bureau, 2004), areas with many housing supply constraints may end up with a smaller share of people in these groups. Furthermore, high house prices may mean that only rich households can afford to move into a location while poorer households are forced out, leading to higher income inequality across locations (Gyourko et al., 2006). An additional effect of changing the composition of the local population may be that the skills of the local workforce will change, thereby altering the industrial composition of local firms as well.

Of course, policies that restrict the expansion of the housing supply are not without benefits. By enforcing quality standards for residential housing units and reducing the negative externalities associated with crowding and undesirable land uses, these regulations can make some local residents better off. The goal of this paper was not to emphasize the negative consequences of housing supply restrictions, but rather to point out that the costs of housing supply regulation will be underestimated if the effects on the labor market are not taken into account. Thus, discussions of land use policy are not complete without considering their implications for labor markets and employment growth in addition to their impact on housing markets.

Appendix A. Constructing an index of housing supply regulation

I create an index of housing supply regulation using information from the following six sources:

1. Wharton Urban Decentralization Project (Wharton) (Linneman et al., 1990). A survey of city planners and local county officials conducted in the late 1980s concerning local development regulations. I combine the answers to the following nine survey questions into a single measure of housing supply regulation: the average length of time for re-zoning permits to be approved, the average length of time for subdivision permits to be approved, the change in zoning approval time from 1983–1988, the fraction of zoning applications approved, a rating of the adequacy of the provision of infrastructure for growth needs, a rating of the importance of regulation and development standards in the development process, the amount of impact fees per housing unit, the presence of growth management policies, and the importance of various methods growth management. The answers to each question are normalized to have a mean of zero and a variance of 1, and then averaged across all nine answers.
2. Regional Council of Governments (RCG) (Segal and Srinivasan, 1985). A survey carried out between 1975 and 1978, which asked members of the Regional Council of Governments to estimate the fraction of suburban land that had been made unavailable for development through government regulations. This measure ranges from 0 (many locations such as Dallas, TX and Rochester, NY) to 43.5 percent (Sacramento, CA).
3. International City Management Association (ICMA). In 1984, the International City Management Association sent out a survey on economic development to municipal chief administrative officers. I use the responses to four questions related to constraints on construction: whether the city requires developers to prepare an environmental impact statement, whether the city is involved in historic preservation, whether the city has altered environmental regulations to facilitate economic development, and whether the city has reformed building or zoning regulations to facilitate economic development. I aggregate the city-level responses to metropolitan-area averages using 1990 land area as weights. The data are available through the FAUI project (see below).[20]
4. Fiscal Austerity and Urban Innovation (FAUI) (Clark and Goetz, 1994). As part of a large survey of local city governments in 1983–1984, chief administrative officers were asked to rate the importance of imposing controls on new construction as a method of limiting population growth on a scale of 1 to 5. The responses for each city are aggregated to metropolitan areas using 1990 land area as weights.
5. National Register of Historic Places (NRHP). The NRHP maintains a comprehensive database of all properties of historic significance in the United States that are listed in the National Register. The database includes historic buildings, sites, districts, structures, and objects and has information on the location, size, date of certification, and other characteristics of each property. I calculate the total amount of land area in each metropolitan area taken up by a historic district or site.[21] Because the information on lot size is generally missing for other types of places, I separately add up the total number of buildings and structures that are not in a district or site. Then, I create a measure of historic preservation equal to the average of the number of buildings and structures per square mile and the fraction land in the metropolitan area that is a part of a historic district or site.
6. American Institute of Planners (AIP) (American Institute of Planners, 1976). The AIP conducted a study of the types of land use planning activity undertaken by each of the 50 states in 1976. I create an index based on the presence or absence of eight different policies with the potential to restrict residen-

[20] I thank Terry Clark for providing the data and original survey documentation.

[21] The database contains information on the county where each property is located, so I aggregate to the 1999 Census definitions of metropolitan areas.

tial construction: comprehensive land use planning, coastal zone management plans, wetlands management regulations, floodplain management, designation of some "critical" locations as requiring state involvement for development, legislation regulating the location and development of new towns, requirements for developers to file environmental impact statements, and regulations preempting local regulations for "developments of greater than local impact." A state is assigned a value of 1 for each type of policy it uses, and the final score for each state is the sum over all eight policies. For metropolitan areas that extend into multiple states, I use a weighted average of the state scores where the weights are determined by the fraction of land that extends into each state.

The comprehensive index of housing supply regulation is a simple average of each of the six sources of information described above. Because the geographic areas covered by the first four surveys are not comprehensive, I impute missing values for each survey from observed values of the other surveys. The imputation for a missing value of survey S in location X is based on an OLS regression of S on all of the other survey measures available in that location. Table A.1 shows the bivariate correlation between each pair of surveys and the number of metropolitan areas that they have in common. Except for the NRHP index, the measures are all reasonably correlated with one another. Because the NRHP index displays a low correlation with all of the other measures, I do not use it in any of the imputation equations. After all missing values are imputed, I rescale each index to have a mean of zero and a standard deviation of one.

Table A.1
Correlation of housing supply survey measures

	Wharton	RCG	ICMA	FAUI	NRHP	AIP
Wharton	1	0.50	0.45	0.60	0.08	0.28
	(60)	(40)	(43)	(27)	(60)	(60)
RCG		1	0.66	0.45	−0.09	0.36
		(48)	(33)	(25)	(48)	(48)
ICMA			1	0.13	0.04	0.06
			(155)	(46)	(155)	(155)
FAUI				1	−0.13	0.02
				(62)	(62)	(62)
NRHP					1	−0.02
					(318)	(318)
AIP						1
						(318)

Notes. Each cell shows the bivariate correlation between each pair of survey responses. The number of observations that each pair has in common is shown in parentheses.

The final index of supply regulation is a simple average of these six components, and I limit the analysis to areas with at least four non-imputed values.[22] The final index covers a total of 82 metropolitan areas and the complete ranking is shown in Table A.2.

A.1. Other variables

Wherever possible, all variables are calculated to reflect the 1999 Census definitions of Primary Metropolitan Statistical Areas and New England County Metropolitan Areas.

Housing stock: Annual values of the housing stock are estimated from annual data on building permits for residential construction and the following stock-flow equation:

$$\text{Stock}_{it} = \text{Stock}_{it-1} + \text{Permits}_{it} - \text{Adjustment factor}_{it}.$$

Permits include both single-family and multi-family units. I begin with the total number of housing units reported in the 1980 Census and calculate decadal adjustment factors so that the 1990 and 2000 values predicted by the stock-flow equation match the Census counts. Estimates of the housing stock in 2001 and 2002 are taken directly from county-level estimates provided by the Census Bureau's Population Estimates.

House prices: The average housing price for each metropolitan area is based on the repeat-transactions price indexes published by the Office of Federal Housing Enterprise Oversight. Because this index calculates price changes from homes that are sold multiple times, changes in the quality of the housing stock will only affect the index when houses are renovated or depreciate significantly. I create annual index values as the annual average of the quarterly values and re-base the index for each location so that its value in 1990 equals the median house price for that location in the 1990 Census. These nominal prices are deflated by the PCE chain-price index. The OFHEO reports indexes based on the 1999 Census definition of PMSAs rather than NECMAs in New England, so for these locations I assign the house price index from the closest possible PMSA to each NECMA.

[22] An alternative method of combining these six surveys would be to extract a common component using factor analysis. This strategy yields a component with an eigenvalue of 2.9 and positive factor loadings on all six surveys. The factor gives about equal weight to each of the four land-use surveys, and a bit less to the index of historic preservation and the state-wide regulations. Thus, results are similar to using the simple average across sources.

Table A.2
Ranking of metropolitan areas by degree of housing supply regulation (high values indicate more regulation)

Index value	Metropolitan area name
−2.40	Bloomington–Normal, IL
−1.96	Buffalo–Niagara Falls, NY
−1.65	Nashville, TN
−1.50	Owensboro, KY
−1.48	Joplin, MO
−1.37	Pueblo, CO
−1.36	Champaign–Urbana, IL
−1.32	Oklahoma City, OK
−1.26	Dayton–Springfield, OH
−1.23	Richmond–Petersburg, VA
−1.18	Dallas, TX
−1.15	Killeen–Temple, TX
−1.07	Lubbock, TX
−1.07	Beaumont–Port Arthur, TX
−1.01	Chicago, IL
−0.95	Kansas City, MO–KS
−0.91	Phoenix–Mesa, AZ
−0.85	Little Rock–North Little Rock, AR
−0.77	Atlanta, GA
−0.69	Detroit, MI
−0.68	Rochester, NY
−0.68	Denver, CO
−0.66	St. Louis, MO–IL
−0.66	San Antonio, TX
−0.65	Grand Rapids–Muskegon–Holland, MI
−0.64	Topeka, KS
−0.55	Indianapolis, IN
−0.52	Houston, TX
−0.47	Greensboro–Winston–Salem–High Point, NC
−0.46	Birmingham, AL
−0.45	Allentown–Bethlehem–Easton, PA
−0.32	Orange County, CA
−0.27	Vallejo–Fairfield–Napa, CA
−0.25	Cleveland–Lorain–Elyria, OH
−0.22	Louisville, KY–IN
−0.20	New Orleans, LA
−0.16	Minneapolis–St. Paul, MN–WI
−0.14	Raleigh–Durham–Chapel Hill, NC
−0.12	Hartford, CT
−0.09	Wichita Falls, TX
−0.07	Columbus, OH
−0.06	Laredo, TX
−0.04	Charlotte–Gastonia–Rock Hill, NC–SC
0.02	Eugene–Springfield, OR
0.08	Sarasota–Bradenton, FL
0.10	Oakland, CA
0.16	Tampa–St. Petersburg–Clearwater, FL
0.16	Cincinnati, OH–KY–IN
0.19	Milwaukee–Waukesha, WI
0.20	Fresno, CA
0.23	Fort Lauderdale, FL
0.26	Pittsburgh, PA
0.27	New Haven–Bridgeport–Stamford–Waterbury–Danbury, CT
0.35	Providence–Warwick–Pawtucket, RI
0.41	Salinas, CA
0.47	Miami, FL
0.47	Philadelphia, PA-NJ
0.48	Austin–San Marcos, TX
0.50	Orlando, FL
0.51	West Palm Beach–Boca Raton, FL
0.56	Albany–Schenectady–Troy, NY
0.59	Melbourne–Titusville–Palm Bay, FL
0.65	Syracuse, NY
0.80	Baltimore, MD
0.86	Boston–Worcester–Lawrence–Lowell–Brockton, MA–NH
0.86	Washington, DC–MD–VA–WV
0.88	Honolulu, HI
0.94	Portland–Vancouver, OR–WA
0.96	Salt Lake City–Ogden, UT
0.98	Tyler, TX
1.02	Newark, NJ
1.15	Ventura, CA
1.21	Los Angeles–Long Beach, CA
1.23	Gary, IN
1.48	Seattle–Bellevue–Everett, WA
1.51	Santa Barbara–Santa Maria–Lompoc, CA
1.60	San Diego, CA
1.65	San Jose, CA
1.73	Riverside–San Bernardino, CA
1.84	Charleston–North Charleston, SC
1.89	Sacramento, CA
2.10	San Francisco, CA
2.21	New York, NY

Employment: Employment is defined as total non-farm employment in each metropolitan area, and is aggregated from county-level data from the Census Bureau's County Business Patterns.

Wages: Wages are calculated as total wage and salary disbursements divided by the number full-time equivalent employees as reported by the Bureau of Economic Analysis. Real values are calculated using the PCE chain-price index.

Unionization: I calculate the average share of unionized workers in each metropolitan area using the Merged Outgoing Rotation Groups of the Current Population Survey from 1986 to 2002. The unionization rate is defined as the number of union members relative to total employment. I exclude the construction industry in order to avoid any issues concerning the effect of union wages on construction costs.

State regulatory index: This index is taken from a state-level study on economic freedom conducted by Byers et al. (1999), and is based on 23 separate measures of labor and school choice regulation, workers compensation, environmental legislation, and regulations in the trucking, insurance, and public utility industries. Most measures pertain to the late 1990s, and no direct measure of housing market regulation is included. Higher values of the index correspond to a stronger degree of regulation.

Metropolitan area age: I classify locations by the decade in which they became a metropolitan area, defined as having a central-city population greater than 50,000 and a total area population greater than 100,000 (Bogue, 1953). Areas are classified into those established before 1900, those established in each decade between 1900 and 1950, and those established after 1950. Metropolitan areas in the oldest category are assigned the highest value so that this variable is increasing in age.

Labor demand shocks: Following Bartik (1991), the predicted change in labor demand is a weighted average of national industry growth rates, where the weights are equal to the share of an industry's employment relative to total metropolitan area employment. To be specific, the formula I use is:

$$\hat{\varepsilon}_{it}^{d} = \sum_{j=1}^{59} \frac{e_{ijt-1}}{e_{it-1}} \left(\frac{\tilde{e}_{ijt} - \tilde{e}_{ijt-1}}{\tilde{e}_{ijt-1}} - \frac{e_t - e_{t-1}}{e_{t-1}} \right)$$

where $i = \text{MSA}$, $j = \text{industry}$, $t = \text{year}$, $\tilde{e}_{ijt} =$ national industry employment outside of MSA $i = e_{jt} - e_{ijt}$, $e_{it} =$ metropolitan area employment $= \sum_{j=1}^{59} e_{ijt}$, $e_t =$ national employment $= \sum_{i=1}^{318} e_{it}$.

The formula above shows that the industry employment growth rates are adjusted to (1) exclude local employment growth when calculating industry employment growth rates and (2) express industry employment growth relative to national employment growth. For the years 1980–1997, employment in each metropolitan area is defined at the 2-digit SIC level using data from the County Business Patterns and aggregated to the metropolitan level using 1999 PMSA and NECMA definitions. For the years 1998–2002, I define industries using 3-digit NAICS. Because the definition of an establishment changes over time in the County Business Patterns data, I use data from the Bureau of Labor Statistics to calculate national industry employment growth rates. I omit the construction industry because changes in construction industry employment are likely to be related to housing supply regulation, and I omit agricultural employment because the BLS does not report employment in these industries.

Low-demand metropolitan areas: I define a metropolitan area as having low housing demand if it experienced population growth of less than 1.1 percent per year and increases in house prices relative to construction costs of less than 0.4 percent per year from 1960 to 1980. These values are roughly the lowest 35th percentile of the distributions of population growth and house price growth. House prices are measured using the median value of single-family homes from the 1960 and 1980 Censuses, where county-level estimates are aggregated to metropolitan areas using the number of single-family homes in 1960 as weights. Construction costs are indexes published by the R.S. Means Company. This methodology identifies the following 27 locations as having low housing demand:

Akron, OH
Albany–Schenectady–Troy, NY
Beaumont–Port Arthur, TX
Binghamton, NY
Buffalo–Niagara Falls, NY
Chicago, IL
Cincinnati, OH–KY–IN
Cleveland–Lorain–Elyria, OH
Dayton–Springfield, OH
Detroit, MI
Flint, MI
Gary, IN
Indianapolis, IN
Jersey City, NJ
Louisville, KY–IN
Mobile, AL
Omaha, NE–IA
Rochester, NY
Rockford, IL
Shreveport–Bossier City, LA
South Bend, IN
Springfield, MA
St. Louis, MO–IL
Syracuse, NY
Toledo, OH
Utica–Rome, NY
Youngstown–Warren, OH

References

American Institute of Planners, 1976. Survey of State Land Use Planning Activity. Department of Housing and Urban Development, Washington, DC.

Barro, R., Sala-i-Martin, X., 1991. Convergence across states and regions. Brookings Papers on Economic Activity 1991 (1), 107–182.

Bartik, T., 1991. Who benefits from state and local economic development policies? W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Blanchard, O., Katz, L., 1992. Regional evolutions. Brookings Papers on Economic Activity 1992 (1), 1–75.

Bogue, D., 1953. Population Growth in Standard Metropolitan Areas 1900–1950. Housing and Home Finance Agency, Washington, DC.

Bover, O., Muellbauer, J., Murphy, A., 1989. Housing, wages and UK labor markets. Oxford Bulletin of Economics and Statistics 51 (2), 97–136.

Byers, J., McCormick R., Yandle B., 1999. Economic freedom in America's 50 states. Manuscript. http://sixmile.clemson.edu/freedom/index.htm.

Case, K., 1991. The real estate cycle and the economy: Consequences of the Massachusetts boom of 1984–97. New England Economic Review (Sept/Oct.), 37–46.

Clark, T., Goetz, E., 1994. The antigrowth machine. In: Clark, T. (Ed.), Urban Innovation: Creative Strategies for Turbulent Times. Sage Publications, Thousand Oaks, CA.

DiPasquale, D., 1999. Why don't we know more about housing supply? Journal of Real Estate Finance and Economics 18, 9–23.

Dubin, J., Kieweit, D., Noussair, C., 1992. Voting on growth control measures: Preferences and strategies. Economics and Politics 4, 191–213.
Evenson, B., 2002. Understanding house price volatility: Measuring and explaining the supply side of metropolitan area housing markets. Working Paper. Illinois State University Department of Economics.
Fischel, W., 1985. The Economics of Zoning Laws. Johns Hopkins Univ. Press, Baltimore, MD.
Gabriel, S., Shack-Marquez, J., Wascher, W., 1993. The effects of regional house price and labor market variability on interregional migration: Evidence from the 1980s. In: Kingsley, G., Turner, M. (Eds.), Housing Markets and Residential Mobility. The Urban Institute Press, Washington, DC.
Gallin, J., 2004. Migration and labor market dynamics. Journal of Labor Economics 22, 1–21.
Glaeser, E., Gyourko, J., 2005. Urban decline and durable housing. Journal of Political Economy 113, 345–375.
Glaeser, E., Gyourko, J., Saks, R., 2006. Urban growth and housing supply. Journal of Economic Geography 6, 71–89.
Green, R., Malpezzi, S., Mayo, S., 1999. Metropolitan-specific estimates of the price elasticity of supply of housing, and their sources. American Economic Review 95, 334–339.
Gyourko, J., Mayer, C., Sinai, T., Superstar cities. Working paper 12355, NBER.
Hamermesh, D., 1993. Labor Demand. Princeton Univ. Press, Princeton, NJ.
Hwang, M., Quigley, J., 2006. Economic fundamentals in local housing markets: Evidence from US metropolitan regions. Journal of Regional Science 46, 425–453.
Johnes, G., Hyclak, T., 1999. House prices and regional labor markets. Annals of Regional Science 33, 33–49.
Linneman, P., Summers, A., Brooks, N., Buist, H., 1990. The state of local growth management. Working paper 81. Wharton Real Estate Center.
Malpezzi, S., 1996. Housing prices, externalities and regulation in US metropolitan areas. Journal of Housing Research 7, 209–241.
Mayer, C., Somerville, C., 2000. Land use regulation and new construction. Regional Science and Urban Economics 30, 639–662.
Saks, R., 2004. Housing supply restrictions across the United States. Wharton Real Estate Review (Fall).
Segal, D., Srinivasan, P., 1985. The impact of suburban growth restrictions on US housing price inflation, 1975–1978. Urban Geography 6, 1–26.
Seidel, S., 1978. Housing Costs and Government Regulations. Center for Urban Policy Research, New Brunswick, NJ.
Topel, R., 1986. Local labor markets. Journal of Political Economy 94, 111–143.
US Bureau of Labor Statistics, 2002. Consumer Expediture Survey. Bureau of Labor Statistics, USA.
US Census Bureau, 2004. Geographic mobility: 2002–2003. Current Population Reports, March.