

Housing Constraints and Spatial Misallocation[†]

By CHANG-TAI HSIEH AND ENRICO MORETTI*

We quantify the amount of spatial misallocation of labor across US cities and its aggregate costs. Misallocation arises because high productivity cities like New York and the San Francisco Bay Area have adopted stringent restrictions to new housing supply, effectively limiting the number of workers who have access to such high productivity. Using a spatial equilibrium model and data from 220 metropolitan areas we find that these constraints lowered aggregate US growth by 36 percent from 1964 to 2009. (JEL E23, J24, J31, R23, R31)

Starting with Hsieh and Klenow (2009), a large number of studies have documented the existence and the costs of factor misallocation across firms. In this paper, we focus on the spatial misallocation of labor across US cities. This analysis is motivated by the observation of a large and growing spatial dispersion of nominal wages across US cities. After conditioning on worker characteristics, the standard deviation of nominal wages (in logs) across US cities in 2009 is twice as large compared to 1964, indicating that labor productivity is increasingly different across American cities. If productivity of labor is vastly different across cities, output can in principle be increased by expanding employment in high productivity cities at the expense of low productivity cities. We argue that the growing dispersion of the nominal wage across cities reflects growing spatial misallocation which ultimately lowers aggregate growth in the United States.

The increase in spatial wage dispersion is driven at least in part by cities like New York, San Francisco, and San Jose, which experienced some of the strongest growth in labor productivity over the last five decades (Moretti 2012). These cities also adopted land use restrictions that significantly constrained the amount of new housing that can be built. As described by Glaeser (2014), since the 1960s coastal US cities have gone through a property rights revolution that has significantly reduced the elasticity of housing supply: “In the 1960s, developers found it easy to do business in much of the country. In the past 25 years, construction has come to face enormous challenges from any local opposition. In some areas it feels as if every neighbor has veto rights over every project.”

*Hsieh: University of Chicago, 5807 S. Woodlawn Avenue, Chicago, IL 60637, and NBER (email: chsieh@chicagobooth.edu); Moretti: University of California, Berkeley, 619 Evans Hall, Berkeley, CA 94720, and NBER (email: moretti@berkeley.edu). Richard Rogerson was coeditor for this article. An earlier version of this paper circulated under the title “Why Do Cities Matter? Local Growth and Aggregate Growth.” We are grateful to Klaus Desmet, Rebecca Diamond, Daniel Fetter, Cecile Gaubert, Ed Glaeser, Erik Hurst, Pat Kline, Steve Redding, Jose Vasquez, and several seminar participants for useful suggestions. We thank Jihoon Sung for research assistance.

[†]Go to <https://doi.org/10.1257/mac.20170388> to visit the article page for additional materials and author disclosure statements or to comment in the online discussion forum.

Misallocation arises because the constraints on housing supply in the most productive US cities effectively limit the number of workers who have access to such high productivity. Instead of increasing local employment, productivity growth in housing-constrained cities primarily pushes up housing prices and nominal wages. The resulting misallocation of workers lowers aggregate output and welfare of workers in all US cities.

This paper measures the aggregate productivity costs of local housing constraints through the prism of a Rosen-Roback model. In a spatial equilibrium, aggregate output and welfare growth are not simply the sum of local shocks in each city. If workers can move across cities, a localized productivity shock in a city affects wages and employment not only in that city but also in other cities. We derive a formula that shows how local shocks aggregate to affect national output and welfare. Aggregate output and welfare growth depend on a weighted average of productivity shocks in each city and the efficiency of the allocation of workers across cities, where the latter depends on the elasticity of housing supply in high productivity cities. If a city with accommodating housing supply experiences productivity growth, local employment rises and workers in other cities benefit from the reallocation of jobs. If instead the city has restrictive housing supply, the reallocation of jobs is limited and productivity growth in the city is dissipated by the higher price of housing.

We use data from 220 metropolitan areas in the United States from 1964 to 2009 to perform two calculations. First, we quantify the effect of spatial misallocation. We find that most of the increased spatial dispersion in the marginal product of labor is due to the growing spatial dispersion in housing prices. In turn, the growing spatial dispersion of housing prices is largely driven by strict zoning laws in cities such as New York and the San Francisco Bay Area with strong productivity growth. We find that the increased spatial misallocation of labor due to housing supply constraints in cities with high productivity growth rates *lowered* aggregate growth between 1964 and 2009 by a significant amount.

In particular, we calculate that increasing housing supply in New York, San Jose, and San Francisco by relaxing land use restrictions to the level of the median US city would increase the growth rate of aggregate output by 36.3 percent. In this scenario, US GDP in 2009 would be 3.7 percent higher, which translates into an additional \$3,685 in average annual earnings.

Second, we calculate the contribution of each US city to aggregate US growth and compare it to an “accounting” measure based solely on the growth of the city’s GDP. The difference reflects the effect of a city’s growth on the efficiency of labor allocation across cities. While the accounting measure suggests that New York, San Francisco, and San Jose’s contribution to aggregate GDP growth between 1964 and 2009 is 12 percent, viewed through the lenses of our model, these cities were only responsible for 5 percent of growth. The difference is because the aggregate benefit of TFP growth in New York and the Bay Area was in part offset by increased misallocation of labor across cities. In contrast, for southern cities the accounting and model-based measures are the same. Due to an elastic supply of housing, much of the growth in the South took the form of employment growth, with no effect on misallocation.

We conclude that local land use regulations that restrict housing supply in dynamic labor markets have important externalities on the rest of the country. Incumbent homeowners in high productivity cities have a private incentive to restrict housing supply. By doing so, these voters de facto limit the number of US workers who have access to the most productive of American cities. In general equilibrium, this lowers income and welfare of all US workers.

This paper builds on four bodies of work. First, we build on the literature on resource misallocation by showing that frictions stemming from the housing market can impede the efficient allocation of resources across cities.¹ Second, our paper builds on the work that uses cities or regions within a country as laboratories to understand differences in income across countries.² We differ from this work by highlighting how the distribution of economic activity across cities itself affects aggregate outcomes in all cities in the country. Third, we build on the research that measures the effect of local housing supply constraints on housing supply and housing prices.³ Our focus is on the *aggregate* impact of such regulations.⁴ Finally, we use a general equilibrium Rosen-Roback model to measure the effect of local land use regulations on aggregate growth.⁵ Other authors have used similar models to measure the effect of state taxes, internal trade frictions, infrastructure, and land misallocation.⁶

The paper is organized as follows. In Section I we present the model. Section II describes the data. Section III discusses how we infer the driving forces in the model from the data. In Section IV we present the main empirical results. Section V discusses extensions of the model. Section VI discusses policy implications.

I. Model

This section presents a Rosen-Roback model where local forces in a city affect local employment, wages, and housing prices. We show how the GDP of the country as a whole and welfare of the workers are determined by aggregating the effects of local forces in all cities.

¹The existing literature on resource misallocation includes papers on labor market frictions (Hopenhayn and Rogerson 1993; Guner, Ventura, and Xu 2008; Gourio and Roys 2014; Garicano, Lelarge, and Van Reenen 2016; and Hsieh and Klenow 2009), financial frictions (Buera, Kaboski, and Shin 2011; Greenwood, Sanchez, and Wang 2013; Midrigan and Xu 2014; and Moll 2014), restrictions in land markets (Restuccia and Santaeulàlia-Llopis 2015), and distortions in output markets (Peters 2016). Restuccia and Rogerson (2013) and Restuccia and Rogerson (2017) provide recent overviews of the literature on misallocation.

²See, for example, Barro and Sala-i-Martin (1992) and Gennaioli et al. (2014).

³Some examples are Mayer and Somerville (2000); Glaeser and Gyourko (2003); Quigley and Raphael (2004); Glaeser, Gyourko, and Saks (2006); Saks (2008); Saiz (2010); Ganong and Shoag (2017); and Diamond (2017).

⁴Hornbeck and Moretti (2018) uses a different approach to estimate the local and aggregate effects of TFP shocks.

⁵There is by now a large literature on systems of cities in general equilibrium. Some examples are Henderson and Ioannides (1981); Henderson (1982); Behrens, Duranton, and Robert-Nicoud (2014); and Eeckhout, Pinheiro, and Schmidheiny (2014).

⁶See Fajgelbaum et al. (2015) for state taxes, Redding and Turner (2014) for internal trade frictions, Duranton et al. (2015) for land misallocation in India, and Ahlfeldt et al. (2015) for infrastructure.

A. Perfect Mobility

In our setting, workers choose the city that maximizes utility. We first consider the case of perfect mobility across cities. In this case, workers have homogeneous tastes and therefore are infinitely willing to relocate to cities where wages net of cost of housing and amenities are higher. In this case, the local labor supply to a city—namely the number of workers who are willing to relocate for a higher wage or better amenities—is infinitely elastic. We later consider how our results generalize in the case of imperfect mobility. In this case, workers care not only about wages and amenities, but also have heterogeneous tastes over locations. Cities with higher wages, lower cost of housing, or better amenities do attract more workers, but not an infinite number of workers.

Local Employment, Wages, and Housing Prices.—A city indexed by i produces a homogeneous good with the following technology:

$$(1) \quad Y_i = A_i L_i^\alpha K_i^\eta T_i^{1-\alpha-\eta}.$$

Here A_i captures the productivity of the city, L_i employment, K_i capital, and T_i land available for business use. This production function makes three assumptions. First, it assumes that all cities produce the same product. We show in Section IV that this model is isomorphic to one where each city produces a differentiated product and imports products made in other cities. Second, it assumes the production function elasticities α and η are the same in all cities. We will also relax this assumption in Section IV. Third, we assume that the production function is constant returns to scale in capital, labor, and land. Since the supply of land T_i is fixed in each city, this production function is isomorphic to one without land with decreasing returns to scale in labor and capital.

We equate the marginal product of labor and capital to the local nominal wage W_i and the cost of capital R , respectively. We assume that the interest rate R is determined exogenously in world capital markets (and the same in all cities). The (inverse) local labor demand is then

$$(2) \quad L_i = \left(\frac{\alpha^{1-\eta} \eta^\eta}{R^\eta} \cdot \frac{A_i}{W_i^{1-\eta}} \right)^{\frac{1}{1-\alpha-\eta}} \cdot T_i.$$

Labor demand is increasing in A_i and T_i and decreasing in W_i . In what follows, we will refer to the composite $A_i^{\frac{1}{1-\alpha-\eta}} T_i$ as “local TFP.”

Labor supply in a city is pinned down by the condition that workers are freely mobile and in equilibrium choose the city that maximizes utility. Indirect utility of workers is given by

$$(3) \quad V = \frac{W_i Z_i}{P_i^\beta},$$

where P_i denotes the local housing price, β the expenditure share on housing, and Z_i the value of local amenities.⁷

Two features are worth highlighting. First, workers have homogeneous tastes over locations. They do care about wages, housing costs, and amenities, but there is no individual specific preference for a city over another. They are willing to relocate for a infinitesimally small difference in wages, cost of living, or amenities. This assumption, combined with the assumption that workers are completely mobile across cities, implies that the labor supply facing a city is infinitely elastic. Second, we assume the wage is the only source of income. We implicitly assume the housing stock and the land used by businesses are owned by an absentee landlord in another country. We relax both assumptions below.

The local housing price is given by

$$(4) \quad P_i = \bar{P}_i L_i^{\gamma_i},$$

where γ_i is the (inverse) elasticity of housing supply with respect to the number of workers in the city and \bar{P}_i denotes the part of the local housing price that does not vary with employment. Note that we allow γ_i to differ across cities. There is abundant evidence that the housing supply elasticity varies significantly across US cities and that it is an important determinant of local housing costs.⁸ Cities with a limited amount of land or stringent land use regulations have a lower elasticity of housing supply (large γ_i), and cities with abundant land or permissive land use regulations have higher elasticity (small γ_i). Increases in the number of workers in a city have a larger effect on housing costs when the elasticity of housing supply is small (γ_i is large).

We can use the definition of indirect utility (3) to express the nominal wage as a function of utility V and the ratio of local housing prices and local amenities:

$$(5) \quad W_i = V \cdot \frac{P_i^\beta}{Z_i} = V \cdot \frac{\bar{P}_i^\beta L_i^{\beta\gamma_i}}{Z_i}.$$

The second equality of equation (5) substitutes equation (4) for the local housing price into the expression for the local wage. Differences in nominal wages across cities reflect differences in city size, local amenities, and the local elasticity of housing supply. The nominal wage is increasing in a city's employment with a city-specific elasticity that depends on the elasticity of housing supply and decreasing the value of local amenities. Conditional on having the same employment and amenities, the nominal wage is lower in a city with a more elastic housing supply.

⁷The share of expenditures on housing does not vary with income in the data (Davis and Ortalo-Magné 2011, Lewbel and Pendakur 2009), which suggests that β is roughly constant. An alternative model would be a utility function where the share of expenditures on housing changes endogenously with the housing price. We do not pursue this alternative in this paper.

⁸See, for example, Glaeser, Gyourko, and Saks (2006); Saks (2008); Glaeser and Ward (2009); and Saiz (2010).

Equilibrium employment in a city is then given by equating labor demand (2) with labor supply (5):

$$(6) \quad L_i = \left(\frac{\alpha^{1-\eta} \eta^\eta}{R^\eta V^{1-\eta}} \cdot A_i T_i^{1-\alpha-\eta} \cdot \left(\frac{Z_i}{\bar{P}_i^\beta} \right)^{1-\eta} \right)^{\frac{1}{1-\alpha-\eta+\beta\gamma_i(1-\eta)}}.$$

Differences in employment across cities are driven by differences in local TFP, amenities, and the elasticity of housing supply. Cities with more employment are those with high local TFP and high quality amenities, with an elasticity that is increasing in the elasticity of the local housing supply.

Aggregate Output and Welfare.—We now impose the condition that aggregate labor demand is equal to aggregate labor supply (normalized to one). This gives us the following expression for aggregate output $Y \equiv \sum_i Y_i$:

$$(7) \quad Y = \left(\frac{\eta}{R} \right)^{\frac{\eta}{1-\eta}} \left[\sum_i \left(A_i \cdot \left[\frac{\bar{Q}}{Q_i} \right]^{1-\eta} \right)^{\frac{1}{1-\alpha-\eta}} \cdot T_i \right]^{\frac{1-\alpha-\eta}{1-\eta}},$$

where $Q_i \equiv P_i^\beta/Z_i$ is the ratio of the housing price P_i to local amenities Z_i , $\bar{Q} \equiv \sum_i L_i Q_i$ is the employment-weighted average of Q_i across all cities, and the housing price is determined by equation (4). In what follows, for brevity we will refer to $Q_i \equiv P_i^\beta/Z_i$ as the “local price,” by which we mean the ratio of housing prices to amenities. Aggregate output is a power mean of local TFP weighted by the inverse of the local price in the city relative to the average in all cities. Importantly, since $(1-\eta)/(1-\alpha-\eta) > 1$, a mean preserving spread of local prices lowers aggregate output. The aggregate effect of the dispersion in local prices is increasing in the labor share α .

Note the similarity of equation (7) with the expression for aggregate TFP in Hsieh and Klenow (2009). In Hsieh and Klenow (2009), aggregate TFP is increasing in the power mean of firm TFP and decreasing in the dispersion of the firm-specific frictions. Here, aggregate output is increasing in the power mean of TFP of individual cities and decreasing in the dispersion of local prices across cities with an elasticity that is increasing in the labor share α .

It is important to note that in equilibrium, local wages and local prices are linked, because workers need to be indifferent across cities. In particular, the (inverse of) the local price in a city Q_i relative to the national average \bar{Q} is equal to the wage in the city relative to the national average:

$$(8) \quad \frac{\bar{Q}}{Q_i} = \frac{W_i}{\bar{W}},$$

where $\bar{W} \equiv \sum_i L_i W_i$ is the employment-weighted average of W_i across all cities.⁹ Since the local nominal wage is equal to the marginal product of labor in a city,

⁹Equation (8) follows immediately from the equilibrium condition that workers are indifferent across all cities so $V \equiv W_i P_i^\beta/Z_i$ is the same everywhere.

dispersion in local prices results in a misallocation of labor across cities in the same way that dispersion in firm-specific frictions in Hsieh and Klenow (2009) generate misallocation of labor across firms. Intuitively, if the marginal product of labor is different across cities, output can be increased by moving labor from cities with low marginal product to cities with high marginal product, until the marginal product is equalized. And this effect is larger when labor is more important (when α is larger).

Equations (7) and (8) make explicit the link between the housing and labor markets and aggregate output. Differences across cities in housing markets and amenities show up as differences in wages, which generate differences across cities in the marginal product of labor. In this sense, this paper is about how differences in housing prices and amenities across cities generate gaps in the marginal product of labor across firms located in different cities and to what extent these gaps cause misallocation. More specifically, our focus is on the role played by the elasticity of housing supply—the parameter γ_i in the model, which varies across cities—in driving spatial variation in housing prices and therefore in marginal product of labor.

Aggregate utility is given by the ratio of aggregate labor income to the average local price across all cities \bar{Q} . Since the labor share of income is α , aggregate utility is given by

$$(9) \quad V = \alpha \cdot \frac{Y}{\bar{Q}},$$

where aggregate output Y is given by equation (7).

The equilibrium of the model is defined by the equations for local employment (2), housing prices (4), the nominal wage (5), aggregate output (7), and aggregate utility (9). The “exogenous” variables are local TFP, local amenities, and the local elasticity of residential housing supply in all US cities. These variables collectively determine local outcomes in each city (housing prices, wages, and employment) and outcomes for the country as a whole (output and welfare). In Section IV, we discuss the case of endogenous local TFP and amenities.

B. Imperfect Labor Mobility

We now relax the assumption of perfect labor mobility. We assume that workers differ in preferences over locations (see, for example, Moretti 2012 and Kline and Moretti 2014). Specifically, suppose the indirect utility of worker j in city i is given by

$$V_{ji} = \epsilon_{ji} \cdot \frac{W_i Z_i}{P_i^\beta},$$

where ϵ_{ji} is a random variable measuring preferences for city i by individual j . A larger ϵ_{ji} means that worker is particularly attached to city j for idiosyncratic reasons. We assume that workers locate in the city where her utility V_{ji} is maximized. The key difference is that only *marginal* workers are indifferent across cities, but most workers are *infra-marginal* in that their utility in their chosen city is higher than in

all other cities. The implication is that the marginal workers will relocate when real wages and amenities change, but there will be some infra-marginal workers that will not.

To make this model tractable, we assume that ϵ_{ji} are independently distributed and drawn from a multivariate extreme value distribution. Specifically, we follow Kline and Moretti (2014) and assume the joint distribution of ϵ_{jt} is given by $F_g(\epsilon_1, \dots, \epsilon_N) = \exp(-\sum_i^N \epsilon_i^{-\theta})$ where the parameter $1/\theta$ governs the strength of idiosyncratic preferences for location and therefore the degree of labor mobility.¹⁰

The key change then is that labor supply is now upward sloping where the slope depends on the heterogeneity of idiosyncratic location preferences. Specifically, the (inverse) local labor supply to a city is given by

$$(10) \quad W_i = V \cdot \frac{P_i^\beta L_i^{1/\theta}}{Z_i},$$

where V now denotes *average* worker utility in all cities. The elasticity of the labor supply curve depends on the strength of location preferences $1/\theta$. Intuitively, if a city experiences an increase in wages or amenities, the number of workers willing to relocate there depends on $1/\theta$. When preferences for location are important, $1/\theta$ is large, the elasticity of labor supply is low, and few workers are willing to move in response to wage or amenity differences. However, if most workers are not attached to their city and will move in response to wage or amenity differences, $1/\theta$ is small, and the elasticity of local labor supply is high. In the extreme case where workers have no heterogeneous preferences for location, $\theta = \infty$ and the elasticity of local labor supply is infinite. This is the perfect mobility case described above.

Equilibrium employment in a city is now given by

$$(11) \quad L_i = \left(\frac{\alpha^{1-\eta} \eta^\eta}{R^\eta V^{1-\eta}} \cdot A_i T_i^{1-\alpha-\eta} \cdot \left(\frac{Z_i}{\bar{P}_i^\beta} \right)^{1-\eta} \right)^{\frac{1}{1-\alpha-\eta+\beta(\gamma_i+1/\theta)(1-\eta)}}.$$

As in the free mobility case, TFP and amenities have a larger effect on local employment when the local housing supply is elastic. But compared to the free mobility case, the elasticity of employment with respect to local TFP and amenities is lower, where the elasticity is decreasing in $1/\theta$. Intuitively, when workers have strong preferences for a location, fewer workers relocate in response to changes in productivity.

Aggregate output is now given by

$$(12) \quad Y = \left(\frac{\eta}{R} \right)^{\frac{\eta}{1-\eta}} \left[\sum_i \left(A_i \cdot \left[\frac{\bar{Q}}{Q_i} \right]^{1-\eta} \cdot T_i^{1-\alpha-\eta} \right)^{\frac{1}{(1-\eta)(1+1/\theta)-\alpha}} \right]^{\frac{(1-\eta)(1+1/\theta)-\alpha}{1-\eta}}$$

and *average* welfare by $V = \alpha \cdot Y/\bar{Q}$, where $\bar{Q} \equiv \sum_i L_i^{1+1/\theta} \cdot Q_i$. The aggregate effect of a mean preserving spread in $Q_i \equiv P_i^\beta/Z_i$ in the imperfect mobility case

¹⁰None of the substantive results here hinge on the extreme value assumption. See Kline (2010) and Busso, Gregory, and Kline (2013) for analyses with a nonparametric distribution of tastes.

depends on $(1 - \eta)/((1 + 1/\theta)(1 - \eta) - \alpha)$. The effect of dispersion in housing prices and amenities is larger when α is large (as before) and when preferences for location are weak ($1/\theta$ is small).

Compared to the free mobility case, the aggregate effect of heterogeneity in housing prices and amenities is smaller. Intuitively, differences in housing prices and amenities have a smaller effect on the local wage because of imperfect mobility. The elasticity of the local wage with respect to housing prices or amenities is $(1 - \alpha - \eta)/[(1 + 1/\theta)(1 - \eta) - \alpha] < 1$. When $\theta = \infty$, then we are in the perfect mobility world. The elasticity of the local wage to housing prices or amenities is then 1 and aggregate output given by (7).

C. Intuition

We now illustrate the intuition for how local forces show up in local employment and wages, and also how they determine aggregate output and welfare. We mostly focus on the perfect mobility case, but the intuition is the same with imperfect worker mobility.

Consider the effect of an increase in TFP in a city. At the local level, this raises local employment, housing prices, and wages, where the elasticity of housing supply γ_i determines whether local employment or housing prices increase. If the housing supply is elastic, the increase in local TFP has a large effect on employment and a small effect on wages and housing prices. In the extreme case where housing supply is perfectly elastic, there is no change in housing prices or wages (relative to other cities). However, if housing supply is perfectly inelastic, there is no change in employment, and the TFP shock is fully reflected in higher housing prices and wages.

At the aggregate level, the increase in local TFP has three effects. First, an increase in TFP in a city raises the weighted average of local TFP. From combining the equations for aggregate output (7) and welfare (9), the magnitude of this effect is given by the change in $A_i^{\frac{1}{1-\alpha-\eta}}$. Intuitively, higher TFP raises how much output is produced in the nation and therefore its welfare. We call this the “direct effect” of an increase in local TFP.

Second, there is a “misallocation effect,” due to the fact that the TFP shock changes the local marginal product of labor relative to the rest of the country (\bar{Q}/Q_i in equations (7) and (8)). If the marginal product of labor in the city is affected by the TFP shock initially above the nationwide mean, the increase in the housing price increases the gap between the marginal product of labor in the city and in other cities, worsening misallocation. This effect offsets the direct benefit of higher local TFP on output and welfare. Importantly, the increase in misallocation is larger when the housing supply is inelastic. (Of course, if the local housing price is initially below the nationwide mean, misallocation declines).

Third, there is a “price effect”: the increase in local TFP raises the local price, and therefore the average of local prices in the country (\bar{Q} in equation (9)). This lowers aggregate welfare (but not output). Intuitively, this is akin to higher goods prices that lower real income (for a given level of nominal income). The magnitude of this effect is larger when the local housing supply is inelastic because the effect on the local price is larger.

Putting together the “direct,” “misallocation,” and “price” effects, the change in welfare due to a change in local TFP is

$$(13) \quad \Delta V \propto \Delta \left(\frac{A_i}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{1-\alpha-\eta}}.$$

The numerator measures the “direct” effect of TFP and the denominator the “misallocation” and “price” effects.

We note that the change in local employment is proportional to¹¹

$$(14) \quad \Delta L_i \propto \Delta \left(\frac{A_i Z_i^{(1-\eta)}}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{1-\alpha-\eta}}.$$

When local amenities Z_i are fixed, the change in employment (equation (14)) is exactly equal to the contribution of TFP growth in a city to aggregate welfare (equation (13)). So a city’s employment is a sufficient statistic of the “direct,” “misallocation,” and “price” effect of local TFP on aggregate welfare growth.

A similar logic holds for the case where local amenities change and TFP is fixed. An improvement in amenities generates “direct,” “misallocation,” and “price” effects, where the elasticity of housing supply determines the magnitude of the “misallocation” and “price” effects.¹² In general, when both TFP *and* amenities of a given city change, the resulting change in aggregate welfare is given by

$$(15) \quad \Delta V \propto \Delta \left(\frac{A_i Z_i^{(1-\eta)}}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{1-\alpha-\eta}}.$$

The numerator measures the “direct” effect of TFP and amenities and the denominator the “misallocation” and “price” effects. This is equal to the change in the city’s employment in equation (14) when TFP and amenities change. In sum, the change in employment is a sufficient statistic for the aggregate effect of *all* the local forces in the model (TFP, amenities, and the endogenous change in housing prices).

This “sufficient” statistic result does not depend on the degree of labor mobility. In the case with imperfect labor mobility, the effect of a TFP change on welfare is given by

$$(16) \quad \Delta V \propto \Delta \left(\frac{A_i}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{(1+1/\theta)(1-\eta)-\alpha}}.$$

As one might expect, the aggregate effect of a TFP change depends on the degree of labor mobility as parameterized by $1/\theta$. However, the effect of local TFP on

¹¹ We substitute the expression for the nominal wage (5) into the labor demand equation (2).

¹² The direct effect of better amenities on aggregate utility is given by $Z_i^{\beta(1-\eta)\frac{1}{1-\alpha-\eta}}$. The misallocation and price effects are given by $P_i^{-\beta(1-\eta)\frac{1}{1-\alpha-\eta}}$. The effect on welfare, taking all three effects into account, is proportional to the change in $\left(\frac{Z_i^{(1-\eta)}}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{1-\alpha-\eta}}$. This is proportional to the change in employment in (14) when TFP is held fixed.

employment also depends on the degree of labor mobility. The change in local employment due to a change in local TFP is

$$(17) \quad \Delta L_i \propto \Delta \left(\frac{A_i Z_i^{(1-\eta)}}{P_i^{\beta(1-\eta)}} \right)^{\frac{1}{(1+1/\theta)(1-\eta)-\alpha}}.$$

So as in the perfect mobility case, when Z_i is fixed, the change in employment is equal to the aggregate effect of local TFP growth. Intuitively, when labor is less mobile, a given change in TFP has a smaller aggregate effect, but it also has a smaller effect on local employment.

Looking forward to the empirical section, we will provide two sets of estimates. First, we will use employment data to measure the contribution of specific cities to aggregate growth. This contribution can come from TFP growth, an improvement in local amenities, or an elastic housing supply, and the change in employment measures the net effect of all these forces.

Second, we will decompose the contribution of TFP, amenities, and housing supply to aggregate growth. For this second calculation, employment data is not enough. We also need to know the parameters of the production function (α and η), the local elasticity of housing supply γ_i , the degree of labor mobility θ , and data on wages and housing prices. This decomposition will be sensitive to the model parameters, but conditional on the observed data on employment, the *net* effect of these forces on aggregate welfare does not depend on model parameters.

II. Data

The main data we use are the 1964, 1965, 2008, and 2009 County Business Patterns (CBP). We supplement this data with the 1960 and 1970 Census of Population, the 2008 and 2009 American Community Survey (ACS), and the 1964 and 2009 Current Population Survey (CPS). Since the earliest year for which we could find city-industry level data on wages and employment is 1964, we focus on changes between 1964 and 2009.¹³

Data on employment and average wages are available at the county and county-industry level from the CBP and are aggregated to MSA and MSA-industry level. The main strength of the CBP is its fine geographical-industry detail and the fact that data are available for as far back as 1964.¹⁴ The main limitations of the CBP are that it does not provide worker-level information but only provides county aggregates, and that it lacks information on worker characteristics. Obviously, differences in worker skill across cities can be an important factor that affects average wages. In addition, union contracts may create a wedge between the marginal product of labor and the wage, as union wages may contain economic rents. We augment the

¹³The Appendix provides additional details about the data. Table A1 in the Appendix presents the summary statistics of the data.

¹⁴The published tabulations of the Census of Population provide MSA-level averages of worker characteristics, but the individual-level data on employment and salary with geocodes are not available from the public version of the Census of Population on a systematic basis until 1980. Only a third of metro areas are identified in the 1970 census.

CBP data with MSA-level information on worker characteristics from the Census of Population, the ACS, and the CPS: three levels of educational attainment (high school dropout, high school, college), race, gender, age, and union status.

To purge the average wage of differences in worker characteristics across cities, we calculate a residual wage that conditions for geographical differences in the composition of the workforce. Specifically, we use nationwide individual-level regression based on the CPS in 1964 and 2009 to estimate the coefficients on worker characteristics, and use those coefficients to compute residual wages based on city averages.¹⁵ We end up with a balanced sample of 220 MSAs in 1964 and 2009.¹⁶

Data on the local elasticity of housing supply are from Saiz (2010). For each MSA, these data provide the overall elasticity of housing supply γ_i as well as its two main determinants: land availability and land use regulations.¹⁷

III. Data Inference and Stylized Facts

We now describe how we use the structure of the model to back out the forcing variables of the urban system from the data. The inference is the same in the free and imperfect mobility models, except for the case when we back out amenities Z_i . We also discuss some important facts about geographical differences in our key variables.

Employment.—Remember that the contribution of a city to aggregate welfare growth can be inferred from the city's employment growth. Table 1 (top panel) presents (log) average employment of four groups of US cities in 1964 and 2009 relative to average employment in all US cities. The four groups are New York, San Jose, and San Francisco; the Rust Belt (37 cities); southern cities (96 cities); and other large cities—defined as cities with 2009 employment above 600,000 and not in the other three groups (19 cities).

Three facts stand out. First, and unsurprisingly, employment in the Rust Belt declined in relative terms between 1964 and 2009. Second, and more surprisingly, relative employment in New York, San Francisco, and San Jose also fell. Despite torrid growth in labor demand—mainly driven by the growth of finance, high tech, biotech, and other industries concentrated there—these three cities are smaller today than they used to be (relative to other cities). Third, employment grew in the South, which includes booming cities such as Houston, Atlanta, and Dallas. As we will see later, these differences have important implications for aggregate growth.

¹⁵The residual wage is defined as $W - X'b$, where W is a vector of the average wage in the MSA, X is the vector of average worker characteristics in the MSA, and b is a vector of coefficients on worker characteristics estimated from individual-level regressions in nationwide samples. The Appendix shows that in 2009 the estimated average residual wage obtained from MSA-level data is highly correlated with the average residual wage obtained from individual level data. We cannot do the same for 1964, which is why we rely on MSA-level data.

¹⁶These MSAs account for 71.6 percent and 72.8 percent of US employment in 1964 and 2009, respectively, and 74.3 percent and 76.3 percent of the US wage bill in 1964 and 2009.

¹⁷The data are not perfect. First, the measure of land unavailability is partial as it ignores publicly-owned land. Second, it does not include any information on relative locations within cities. Third, the data ignore the kink in the housing supply curve (Glaeser and Gyourko 2005).

TABLE 1—EMPLOYMENT, AVERAGE WAGES, TFP, HOUSING PRICES, AND AMENITIES

	1964	2009
<i>log employment</i>		
New York, San Francisco, San Jose	2.89	2.55
Rust Belt cities	1.63	0.96
Southern cities	0.82	1.14
Other large cities	2.68	2.23
<i>log residual wage</i>		
New York, San Francisco, San Jose	0.041	0.465
Rust Belt cities	0.072	-0.121
Southern cities	-0.038	-0.037
Other large cities	0.010	0.046
<i>log TFP</i>		
New York, San Francisco, San Jose	3.81	7.14
Rust Belt cities	2.77	1.14
Southern cities	1.14	1.95
Other large cities	3.36	3.68
<i>log housing price</i>		
New York, San Francisco, San Jose	0.409	0.610
Rust Belt cities	0.125	-0.104
Southern cities	-0.128	0.106
Other large cities	0.225	0.333
<i>log amenities</i>		
New York, San Francisco, San Jose	0.094	-0.174
Rust Belt cities	-0.040	-0.049
Southern cities	-0.065	-0.026
Other large cities	0.034	0.020

Notes: The sample includes 220 metropolitan areas observed in both 1964 and 2009. There are 37 Rust Belt cities, 86 southern cities, and 19 other large cities. The table presents the employment-weighted average of each group of cities relative to the weighted average in all 220 cities in the year. Residual wage controls for educational attainment (high school dropout, high school, college), race, gender, age, and union status in each metropolitan area.

Figure 1 plots the distribution of log employment for all the cities in our sample. There is a small decrease in the dispersion of city size, mostly driven by the disappearance of the right and left tails in 1964. The right tail in 1964 includes metropolitan areas such as Chicago, Los Angeles, New York, Boston, Philadelphia, Detroit, Cleveland, Pittsburgh, St. Louis, and Baltimore. Employment in all these cities shrunk between 1964 and 2009. The left tail in 1964 includes cities such as Fort Walton Beach, Florida; Bremerton, Washington; and Fort Collins, Colorado. Employment in the 20 smallest US cities in 1964 increased by an average of 0.81 log points between 1964 and 2009. At the same time, relative employment in the southern cities increased between 1964 and 2009. Among these cities, Atlanta; Houston; Washington, DC; and Dallas joined the group of the ten largest cities by 2009.

Wages.—The amount of spatial misallocation depends on geographical differences in wages. In particular, remember from equations (7) and (8) that misallocation depends on the dispersion across cities in local prices, which in equilibrium is simply the dispersion in the nominal wage. Figure 2 plots the employment-weighted distribution of the nominal wage across cities in 1964 and 2009. Panel A of Figure 2




FIGURE 1. SPATIAL DISTRIBUTION OF EMPLOYMENT

Notes: The graph shows the distribution of de-meaned log employment across MSAs weighted by MSA employment in the relevant year. The sample includes 220 metropolitan areas observed in both 1964 and 2009.

presents the *unconditional* average wage, and panel B presents the *residual* average wage after controlling for differences in worker characteristics across cities. The standard deviation of the unconditional wage across US cities increases by 0.073 log points between 1964 and 2009. After conditioning on observables, the standard deviation of the residual wage increases by even *more* (0.083 log points) over the same period.

Figure 2 also shows that the right tail of the wage distribution in 2009 is significantly thicker than in 1964. The bump in the right tail in the wage distribution in 2009 includes New York, San Francisco, and San Jose. Table 1 (second panel) shows that the average residual wage (relative to the mean residual wage) in New York, San Francisco, and San Jose increased from 0.04 to 0.4 log points between 1964 and 2009. The left tail in 2009 is thicker than in 1964 and the mass of the distribution in 1964 where wages are slightly above the mean wage has hollowed out. This change is driven in part by the cities in the Rust Belt, where the mean residual wage (relative to other cities) fell by almost 0.2 log points between 1964 and 2009.

It is of course possible that heterogeneity in unobserved worker quality across cities also increased between 1964 and 2009. We cannot entirely rule this out, but remember that the increase in the dispersion in the residual wage is even larger than the increase in the dispersion of the unconditional wage. Also, using data from the National Longitudinal Survey of Youth, we find that workers in high wage cities tend to have higher Armed Forces Qualification Test scores, but the correlation disappears when we control for education, race, and ethnicity.¹⁸

¹⁸Recent evidence based on longitudinal data that follow workers moving from low wage cities to high wage cities also suggests that unobservable differences in skill are limited once education is controlled for. Baum-Snow and Pavan (2012) finds that sorting on unobserved ability within education group contributes little to observed




FIGURE 2. SPATIAL DISTRIBUTION OF NOMINAL WAGES

Notes: The graphs show the distribution of de-meaned log wages across MSAs weighted by MSA employment in the relevant year. Conditional wage controls for three levels of educational attainment (high school dropout, high school, college), race, gender, age, and union status in each metropolitan area. The sample includes 220 metropolitan areas observed in both 1964 and 2009.

TFP.—Since local employment is a function of local TFP and the local wage, we can invert this relationship to express local TFP as a function of employment and wages. Specifically, equation (2) can be expressed as

$$A_i^{\frac{1}{1-\alpha-\eta}} \cdot T_i \propto L_i \cdot W_i^{\frac{1-\eta}{1-\alpha-\eta}}.$$

differences in wages across cities of different size. De La Roca and Puga (2017) finds that workers in larger and higher wage cities do not have higher unobserved initial ability, as measured by the individual fixed effects in a wage regression. These findings are consistent with Glaeser and Maré (2001).




FIGURE 3. SPATIAL DISTRIBUTION OF LOCAL TFP

Notes: Local TFP is defined as $A_i^{1-\alpha-\eta}T_i$. The graph shows the distribution of de-meaned log local TFP weighted by MSA employment in the relevant year. The sample includes 220 metropolitan areas observed in both 1964 and 2009.

The left-hand side of this equation is “local TFP” and the right-hand side is a function of local employment, wages, and the parameters η and α .¹⁹ For our baseline, we follow Bureau of Economic Analysis (2013), Piketty (2014), and Karabarbounis and Neiman (2014) and assume $\alpha = 0.65$ and $\eta = 0.25$, which imply a residual share $1 - \alpha - \eta$ of 10 percent.²⁰ Later, we probe the robustness of our results to alternative values of the parameters (Section V). Although our estimates of TFP are model-driven, they match well independent estimates of TFP obtained from directly estimating production functions. In particular, the cross-sectional correlation between our 2009 TFP estimates and TFP estimates obtained by Hornbeck and Moretti (2018) using establishment-level data from the Census of Manufacturers is 0.81.

Figure 3 plots the distribution of local TFP we infer from data on employment and wages. The dispersion of local TFP widened considerably from 1964 to 2009. The employment-weighted standard deviation of local TFP doubled between 1964 and 2009 (from 0.74 to 1.4 log points). This is driven by two changes. First, the mass of the distribution of TFP in 1964 (where TFP was roughly 2 to 3 log points above the mean) has hollowed out. Table 1 indicates that TFP in the Rust Belt cities fell from an average of 2.77 to 1.14 log points above the TFP of the mean US city from 1964 to 2009. Detroit was the city with the highest TFP in 1964. By 2009, TFP in Detroit was below that of the average US city. Second, while

¹⁹Intuitively, the elasticity of employment to the nominal wage depends on η and γ , and the nominal wage captures the effect of housing prices and amenities.

²⁰The assumption that the residual of costs is 10 percent is consistent with Basu and Fernald (1997)’s estimates of the returns to scale in the United States.




FIGURE 4. SPATIAL DISTRIBUTION OF HOUSING COSTS

Notes: The graph shows the distribution of de-meaned log housing costs across MSAs weighted by MSA employment in the relevant year. The sample includes 220 metropolitan areas observed in both 1964 and 2009.

there were no cities in 1964 where relative TFP exceeded 5 log points, this was no longer true in 2009. The top three cities in 2009 are New York, San Jose, and San Francisco. In these three cities, relative TFP increased from 3.8 log points above the mean US city in 1964 to 7.1 log points above the mean city in 2009. This makes sense because these cities are where the most innovative parts of industries such as high tech, biotech, and finance have become increasingly concentrated (Moretti 2012).

It is important to highlight that while the right tail of the distribution of local TFP has thickened, the right tail of the distribution of employment across cities has not. Intuitively, this is because some of the cities with largest TFP gains—most importantly New York, San Jose, and San Francisco—are also some of the most supply-constrained housing markets in the United States. As we will see in detail later, strong TFP gains in these cities have not translated in larger employment but in higher housing costs and a higher marginal product of labor relative to other US cities.

Housing Prices.—Figure 4 plots the distribution of log housing prices. The change in the distribution of housing prices across cities is striking. First, the dispersion has widened: the standard deviation (in logs) increased from 0.11 in 1964 to 0.18 in 2009. Second, the shape of the distribution also changed. In 1964, the mass of the distribution was in cities where housing prices were roughly 0.2 log points higher than prices in the average city. Furthermore, there were no cities where the housing price exceeded the average by more than 0.5 log points. By 2009, however, there was significantly less density in the bins where relative housing prices were roughly 0.2 log points, and significantly more mass where relative housing prices exceed 0.5 log points.




FIGURE 5. SPATIAL DISTRIBUTION OF AMENITIES

Notes: The graph shows the distribution of de-meaned log amenities across MSAs weighted by MSA employment in the relevant year. The sample includes 220 metropolitan areas observed in both 1964 and 2009.

Amenities.—To impute local amenities, we need to take a stance on the degree of labor mobility. When labor is completely mobile across cities, local amenities can be imputed as the residual of the local wage after controlling for the housing price (equation (3)): $Z_i = P_i^\beta / W_i$.²¹ We note that this measure of local amenities is standard in urban economics and has been the subject of much research. Among other things, it has been shown to be highly correlated with local amenities that can be measured like weather, crime, school quality, number of restaurants, and various indices of the quality of life (see, for example, Albouy 2008).

Figure 5 plots the distribution of amenities in the perfect mobility case. The overall dispersion of amenities has not increased very much. The employment-weighted standard deviation is 0.076 in 1964 and 0.1 in 2009. The 1964 distribution is slightly left skewed with a mass at roughly 0.05 above the mean. The 2009 distribution is right skewed with a mass at roughly -0.05 .

With imperfect mobility, the expression for amenities includes an additional term, as explained in Section IB: $Z_i = P_i^\beta L_i^{1/\theta} / W_i$. Recall that the parameter θ governs the degree of worker mobility—the larger the parameter θ , the higher worker mobility.

We calibrate θ based on estimates in Hornbeck and Moretti (2018), which is the empirical study closest to our setting that we are aware of. They estimate the equation of local labor supply (identical to our equation (10)) using arguably exogenous instrumental variables and uncover a long run elasticity of local labor supply

²¹ Since V is unobserved, the absolute level of amenities is not identified. However, the relative level in each city is identified. We use Albouy's (2008) number for the share of expenditures on housing of $\beta = 0.32$. Following Albouy (2008), we also multiply wages by 0.52 to account for taxes and transfers.

$1/\theta = 0.3$. This estimate implies a degree of labor mobility that, while not infinite, is quite high, at least in the long run.²²

As far as amenities are concerned, results under imperfect mobility are qualitatively similar to those under perfect mobility. In particular, the overall dispersion of amenities has not increased and rather has slightly decreased. The employment-weighted standard deviation is 0.68 in 1964 and 0.51 in 2009.

IV. Decomposing Aggregate Growth

Our descriptive evidence in the previous section has uncovered large and growing differences across cities in the marginal product of labor, associated with large and growing differences in housing costs. Consistent with growing misallocation, it appears that some of the cities that have experienced the largest TFP gains—in particular New York, San Francisco, and San Jose—have not expanded their workforce. Rather, they have experienced increases in local costs and wages. We now turn to our empirical estimates of the effect of spatial misallocation. Our model gives us an expression for how local shocks in a city—whether changes in local TFP, amenities, or housing supply—affect aggregate growth. We use our model to decompose aggregate output and welfare growth in order to quantify spatial misallocation and its sources.

We proceed in three steps. First, we estimate the aggregate effects of local growth by quantifying the contribution of specific cities to aggregate growth. The sufficient statistic for the contribution of a city to aggregate growth is the city's employment growth and does not depend on the degree of labor mobility. A city's employment growth captures the aggregate effect of any local shock to a specific city—TFP, amenities, or housing supply—on aggregate output and welfare. We compare model driven estimates to “accounting” estimates based only on the change in local GDP. The accounting estimates differ from the model driven estimates because they ignore the effect of a city's growth on the efficiency of resource allocation across cities.

Having found a large difference, in our second step, we then isolate the contribution of changes in housing prices and amenities on aggregate growth. Having found that housing prices play a major role, in our third step, we then focus on the role of housing supply: we isolate the effect of the local elasticity of housing supply on aggregate growth. These last two calculations depend on the assumption of mobility so we provide two sets of estimates. We first show estimates under the assumption of perfect mobility, then we show how our main estimates change under the assumption of imperfect mobility.

²²Serrato and Zidar (2016) and Diamond (2016) estimate θ to be smaller. We note, however, that both Serrato and Zidar's and Diamond's parameters are unlikely to be the one relevant in our setting, as they are obtained using ten year changes or less. By contrast, Hornbeck and Moretti (2018) uses a 30-year time horizon. A longer time horizon would likely imply more mobility and a larger θ . In this paper, our time horizon is 45 years.

TABLE 2—CONTRIBUTION OF CITIES TO AGGREGATE GROWTH

	Output (percent)	Output (accounting) (percent)	Welfare (percent)
Adjusted aggregate growth rate	0.795	0.795	0.795
<i>Contribution to aggregate growth</i>			
New York, San Francisco, San Jose	5.0	12.3	0.8
Rust Belt cities	15.6	11.1	18.5
Southern cities	32.9	31.9	35.5
Other large cities	33.6	32.4	30.3

Notes: The sample includes 220 metropolitan areas observed in both 1964 and 2009. Row 1 presents average annual adjusted growth rate of aggregate output or welfare between 1964 and 2009. Entries in other rows show the percentage contribution of four groups of cities to aggregate output or welfare growth. There are 37 Rust Belt cities, 86 southern cities, and 19 other large cities. Column 1 presents the model-based estimates for output and column 2 the accounting estimates for output. Column 3 presents the model-based estimates for welfare. See text for details.

A. Aggregate Effect of Local Growth

Table 2 presents the contribution of our four groups of cities to growth in aggregate output (columns 1 and 2) and welfare (column 3). Specifically, we use our model to calculate the net effect of changes affecting a group of cities on aggregate output and welfare. In practice, this calculation assumes that the exogenous forces in our model—local TFP, local amenities, and elasticity of housing supply—of a given group of cities change as in the data while holding these forces in the other cities fixed. We allow employment, housing prices, and wages in all cities to change endogenously until the spatial equilibrium is restored—so that the marginal product of labor is equal to the nominal wage and all workers are indifferent across cities.

Output.—We start by focusing on estimates of the output effects in columns 1 and 2. The first row in Table 2 shows the growth rate of aggregate output observed in the data, which we will use as our benchmark. For this purpose, the relevant growth rate of output is not the growth rate of raw output, but the growth rate of output conditioning on inputs. To measure it, we use the growth of *residual* real per capita earnings.²³ Based on this definition, the adjusted growth rate of aggregate GDP per worker from 1964 to 2009 was 0.795 percent per year.

The second row shows the percentage contribution of New York, San Francisco, and San Jose to aggregate output growth. Column 1 shows that based on our model, these three cities account for 5 percent of aggregate GDP growth. Column 2 shows the contribution of these cities to GDP growth from a naïve accounting calculation

²³Since worker quality increases in this period due to increases in human capital, we used residual earnings. Residual earnings control for differences in education (and other worker characteristics). We note that there are two limitations of our measure. First, labor share has declined. In addition, measured earning growth does not include non-monetary compensation. Since health care and retirement have increased in this period, this second data limitation would lead us to underestimate actual earning growth—and therefore actual output growth.

based on the change in local GDP of these cities as a share of aggregate GDP.²⁴ This calculation suggests that New York, San Jose, and San Francisco were responsible for 12 percent of aggregate growth. The difference between the model-based estimate in column 1 and the “accounting” estimate is that the “accounting” estimate only reflects the “direct effect” of local shocks, while the model-based estimate also takes into account the “misallocation effect.”

The third row in Table 2 shows the contribution of the Rust Belt cities. These cities account for 15.6 percent of US GDP growth. The contribution of the Rust Belt cities to aggregate growth is significantly higher than the contribution of New York, San Jose, and San Francisco. This may seem surprising given the standard narrative of economic decline in the Rust Belt cities and economic dynamism in New York and the Bay Area. The accounting-based estimate of the contribution of the Rust Belt (11 percent of GDP growth) to aggregate GDP growth is lower than the contribution of New York, San Jose, and San Francisco. However, what the accounting estimates miss are the general equilibrium implications of the decline in relative wages in the Rust Belt and the rise in relative wages in New York and the Bay Area. The decline in local prices in the Rust Belt narrows the gap in the marginal product of labor between the Rust Belt and other US cities, while the rise in local prices in New York and the Bay Area increases the gap in the marginal product. This effect is seen in the gap between the accounting estimates and the model-based estimates of the contribution of the Rust Belt cities to aggregate growth (11 percent versus 15.6 percent).

The fourth row in Table 2 shows that the cities in the South account for a large fraction of US growth. In the period under consideration, employment in the US South has grown 56 percent more rapidly than the rest of the country. Examples include Austin (employment growth of +1,002 percent), Washington, DC (+550 percent), Raleigh-Durham (+506 percent), Houston (+363 percent), and Atlanta (+376 percent). The implication of high employment growth is that southern cities account for a large share of aggregate growth. Specifically, southern cities were responsible for 32.9 percent of aggregate GDP growth. Also, the model-based estimate slightly exceeds the accounting estimate, which indicates that growth in these cities improved the allocation of labor across cities.

The last row in Table 2 indicates that the other large cities account for about one-third of aggregate GDP growth. The model-based estimate is about the same as the accounting estimate, suggesting that growth in these cities did not change the efficiency of the allocation of labor across US cities.

Welfare.—We now turn to estimates of the welfare effects. These estimates are shown in column 3 of Table 2. Unlike output growth, which is observed in the data, there is no empirical counterpart of welfare growth. Here we assume the growth rate of aggregate welfare is the same as the growth rate of GDP (row 1). Note that the estimates in Table 2 do not depend on this assumption, because the contribution of

²⁴The accounting decomposition of GDP growth is given by $\frac{Y_t}{Y_{t-1}} = \sum_i \frac{Y_{it}}{Y_{i,t-1}} \cdot \frac{Y_{i,t-1}}{Y_t}$, where Y_t denotes aggregate GDP and $Y_{i,t}$ denotes local GDP of city i (both at time t).

TABLE 3—AGGREGATE EFFECT OF LOCAL CHANGES IN AMENITIES AND HOUSING PRICES

	Output (percent)	Welfare (percent)	Welfare (percent)
Adjusted aggregate growth rate	0.795	0.795	1.795
<i>Percent difference in growth</i>			
Hold housing prices/amenities fixed	112.7	115.0	124.0
Hold amenities fixed	13.6	6.2	6.9
Hold housing prices fixed in			
All cities	103.5	86.8	93.6
New York, San Francisco, San Jose	101.6	73.7	79.6
Rust Belt	4.8	−10.6	−11.4
South	−2.7	18.7	20.2
Other large cities	4.7	15.9	17.1

Notes: The sample includes 220 metropolitan areas observed in both 1964 and 2009. Row 1 presents average annual adjusted growth rate of aggregate output or welfare between 1964 and 2009. Rows 2–4 present the percentage difference in the growth rate when the distribution of P_i^3/Z_i is fixed (row 2), only Z_i is held fixed (row 3), and only P_i is fixed (row 4). Column 2 presents the baseline welfare effect. Column 3 assumes that the rate of growth of welfare between 1964 and 2009 is 1 percent higher. There are 37 Rust Belt cities, 86 southern cities, and 19 other large cities. See text for details.

a given city to welfare growth is the same regardless of the growth rate of aggregate welfare. (Remember that the summary statistic of the relative contribution of a city to welfare growth is the growth of employment.)

The second row in Table 2 shows the percentage contribution of New York, San Francisco, and San Jose to aggregate welfare growth. Their contribution to aggregate welfare growth is 0.8 percent, lower than their contribution to output growth. Remember from Table 1 that the average wage in these three cities increased from 1964 to 2009 relative to wage in the average US city (from 0.041 to 0.465 log points). Viewed through the lenses of the model, this fact indicates that either the housing prices increased or amenities worsened by 0.45 log points in these three cities (relative to the other cities). Higher housing prices or worse amenities in these cities have a negative effect on welfare for a given level of aggregate output (see equation (9)).²⁵

The third row shows the contribution of the Rust Belt cities. These cities account for 18.5 percent of US welfare growth. Their contribution to welfare growth is higher than their contribution to output growth because local prices in the Rust Belt cities declined relative to other cities (as we saw in Table 1). Put differently, the gap in the model-based estimates of the Rust Belt's contribution to GDP (15.6 percent) versus the contribution to welfare (18.5 percent) reflects the effect of the Rust Belt in lowering the average of local prices across all US cities.

The fourth row shows the contribution of cities in the South to welfare growth. Southern cities were responsible for 35.5 percent of welfare growth in the United States. The last row in Table 2 indicates that the other large cities account for about one-third of aggregate welfare growth.

²⁵The difference between the model-based welfare estimate in column 3 and the “accounting” estimate are the “misallocation” and “price” effects.

B. Aggregate Effects of Local Changes in Housing Prices and Amenities

In Table 2, we have shown the net effect of changes in local TFP, amenities, and housing prices in a given city (or group of cities). We now isolate the contribution of amenities and housing prices in Table 3. We do this for the perfect mobility case. Remember that in this case, the equilibrium nominal wage is proportional to P_i^β/Z_i and a localized change in amenities and housing prices in a city affects aggregate output through its effect on the dispersion of nominal wages across cities. When the dispersion of nominal wages across cities increases, the resulting misallocation of labor across cities lowers aggregate output.

Output.—In column 1 of Table 3, we calculate a counterfactual where we hold the distribution of P_i^β/Z_i in all US cities fixed at its 1964 level and assume that local TFP changes as in the data. The entry in row 2 shows the percent difference between the growth rate in this counterfactual and that observed in the data (shown in row 1). The difference between the counterfactual and the actual growth rate of aggregate output is more than a factor of two. This difference is the aggregate effect of the increase in the spatial misallocation of labor across US cities between 1964 and 2009 due to the change in the spatial dispersion of P_i^β/Z_i .

By definition, variation in local prices comes from housing costs P_i and local amenities Z_i . We now isolate the effect of each one at a time. The third row isolates the effect of changes in the dispersion in local amenities. The effect of holding the distribution of local amenities fixed at their 1964 levels appears quantitatively small.²⁶

In the fourth row of Table 3, we compute the counterfactual when we keep the distribution of housing prices fixed at their 1964 levels and allow all the other variables (amenities and local TFP) to vary as in the data. The growth rate of aggregate output under this counterfactual is about twice as large as the actual growth rate, and only slightly lower than the growth rate when we keep the distribution of the ratio of housing prices to amenities fixed.

In sum, changes in the spatial dispersion in housing costs are an important determinant of output growth in the United States in this period, while the aggregate effect of changes in the spatial distribution of amenities is limited. Put differently, the change in the distribution of P_i^β/Z_i (and therefore marginal product of labor) is mostly driven by changes in the distribution of housing prices across cities.

The rest of the table replicates the same exercise for the four groups of cities. Specifically, we hold fixed the gap in housing prices between each group of cities and the average housing price (in all cities) to its 1964 level and calculate the equilibrium under this scenario. In row 5, we show the effect of holding fixed the gap in housing prices only in New York, San Francisco, and San Jose while allowing other local forces (housing prices in the other cities and amenities and local TFP in *all* cities) to change as in the data. We find that the aggregate growth rate more than doubles. Put differently, the increase in housing prices in only these three cities

²⁶For example, column 1 says that if the distribution of amenities across cities had not changed between 1964 and 2009, the growth rate of aggregate output in the United States would have been 13.6 percent higher.

accounts for most of the aggregate output effects. In contrast, when we hold housing prices fixed in the Rust Belt, the South, or in the other large cities, the effect on growth of output is not very large.

Welfare.—The welfare results are shown in columns 2 and 3 of Table 3. (Recall that the welfare effect differs from the output effect only because of changes in the average local price $\bar{Q} \equiv \sum_i L_i P_i^\beta / Z_i$.) As we mentioned above, we do not know the true rate of welfare growth. But for a given rate of growth, we can compute the hypothetical rate under the scenario where the distribution of P_i^β / Z_i is unchanged. Entries in column 2 and 3 are based on the assumption that the rate of growth of welfare is the same as the rate of growth of output or 1 percent higher, respectively.

The empirical results are not sensitive. In both cases, we find that the hypothetical growth rate of aggregate welfare under the scenario where the distribution of P_i^β / Z_i is unchanged differs from the assumed growth rate by more than a factor of two (row 2, columns 2 and 3).

The third row in Table 3 shows that the welfare effect of holding the distribution of local amenities fixed at their 1964 levels appear quantitatively small. In the fourth row, we compute the counterfactual when we keep the distribution of housing prices fixed at their 1964 levels. The growth rate of aggregate welfare under this counterfactual is about twice as large as the actual growth rate, and only slightly lower than the growth rate when we keep the distribution of the ratio of housing prices to amenities fixed.

In row 5 in Table 3, we find that holding fixed the gap in housing prices only in New York, San Francisco, and San Jose would result in a large effect on welfare growth. When compared with the corresponding entries in row 4, it suggests that the increase in housing prices in these three cities accounts for most of the aggregate welfare effects.

In contrast, when we hold housing prices fixed in the Rust Belt (row 6), the growth rate of welfare is *negative*. Glaeser, Gyourko, and Saks (2006) shows that housing prices fall significantly in cities that experience large declines in labor demand, and the experience of the Rust Belt cities is a prominent example of their finding. What Table 3 highlights is that the drop in housing prices has a large *positive* effect on welfare (for a given level of aggregate output). Finally, when we hold housing prices fixed in the South and in the other large cities, the counterfactual growth of welfare is larger than the growth rate of GDP. This indicates that housing prices have increased in these last two groups of cities, but the magnitude is much lower than in New York and the Bay Area.

C. Aggregate Effect of Local Changes in Housing Supply Elasticity

We have shown that the increase in the spatial distribution of housing prices—and in particular changes in prices in New York and the Bay Area—had important negative effects on aggregate output and welfare growth. We now explicitly focus on the role played by local land use restrictions in driving changes in local housing prices and ultimately aggregate growth.

Land use regulations are typically measured using the Wharton Residential Land Use Regulatory Index. This index is based on a detailed survey of municipalities in

TABLE 4—AGGREGATE EFFECT OF LOCAL CHANGES IN HOUSING SUPPLY ELASTICITIES:
PERFECT MOBILITY

	Output (percent)	Welfare (percent)	Welfare (percent)
Adjusted aggregate growth rate	0.795	0.795	1.795
<i>Changing land use regulations in</i>			
New York, San Francisco, San Jose	86.1	51.8	55.8
Rust Belt	0.1	-14.7	-15.5
South	-25.0	-44.7	-48.6
Other large cities	-37.0	-38.3	-41.4

Notes: Row 1 presents average annual adjusted growth rate of aggregate output or welfare between 1964 and 2009. Entries in other rows are estimates of the effect of changing land use regulations in selected cities so that housing supply elasticity in those cities equals the elasticity of the median city. The sample includes 220 metropolitan areas observed in both 1964 and 2009. There are 37 Rust Belt cities, 86 southern cities, and 19 other large cities. Column 2 presents the baseline welfare effect. Column 3 assumes that the rate of growth of welfare between 1964 and 2009 is 1 percent higher. See text for details.

2007. According to the Wharton Index, the restrictions on land use in New York, San Francisco, and San Jose are among the tightest in the country. The elasticity of housing supply due to land use regulations in San Francisco is at the ninety-ninth percentile and New York and San Jose at the ninety-sixth percentile of the nationwide distribution of the housing supply elasticity. In contrast, the average restriction on land use in the US South are in the forty-sixth percentile.

We cannot measure land use restrictions in 1964 because no systematic data exist on land use restrictions in the past (the Wharton Index does not go back in time). We thus illustrate the aggregate effect of land use restrictions by examining the counterfactual where we assume land use regulations in each of our four groups of cities are equal to the level of regulations in the median US city. More specifically, we hold the amount of available land equal to its actual amount, and only change land use regulations to the level observed in the median US city.

In practice, we proceed in three steps. First, we use Saiz's (2010) estimates to compute what would the elasticity of housing supply be in a given city if that city changed its land use regulations to be equal to the median city, holding constant land and everything else.²⁷ Second, with this value of the housing supply elasticity, we estimate the counterfactual equilibrium levels of housing prices, wages, and employment in each group of cities assuming that all the other variables (amenities and local TFP) vary as in the data. Finally, we calculate the counterfactual growth rate of output and welfare with this new allocation of labor.

Output.—Output estimates are in column 1 of Table 4 and Table 5. Starting with perfect mobility, the second row in Table 4 shows the effect of changing the housing supply regulation only in New York, San Jose, and San Francisco to that in the median US city. This would increase the growth rate of aggregate output from 0.795 percent to 1.49 percent per year—an 87 percent increase (column 1). The net effect

²⁷ We use the estimates in Table 5, column 2, in Saiz (2010).

TABLE 5—AGGREGATE EFFECT OF LOCAL CHANGES IN HOUSING SUPPLY ELASTICITIES:
IMPERFECT MOBILITY

	Output (percent)	Welfare (percent)	Welfare (percent)
Adjusted aggregate growth rate	0.795	0.795	1.795
<i>Changing land use regulations in</i>			
New York, San Francisco, San Jose	36.3	13.1	14.1
Rust Belt	2.4	-13.2	-14.2
South	-17.7	-40.8	-44.1
Other large cities	-19.0	-29.5	-31.8

Notes: Row 1 presents average annual adjusted growth rate of aggregate output or welfare between 1964 and 2009. Entries in other rows are estimates of the effect of changing land use regulations in selected cities so that housing supply elasticity in those cities equals the elasticity of the median city. The sample includes 220 metropolitan areas observed in both 1964 and 2009. There are 37 Rust Belt cities, 86 southern cities, and 19 other large cities. Column 2 presents the baseline welfare effect. Column 3 assumes that the rate of growth of welfare between 1964 and 2009 is 1 percent higher. See text for details.

is that US GDP in 2009 would be 8.9 percent higher under this counterfactual, which translates into an additional \$8,775 in average wages for all workers.²⁸ In this counterfactual, wages in New York, San Jose, and San Francisco would be on average 25 percent lower and employment would be higher. Intuitively, the housing supply is more accommodating in this counterfactual. Therefore, more workers can move to these three cities from the rest of the United States.

Housing supply is generally rather elastic in southern cities. This reflects abundant land and permissive land use regulations in southern cities. To see the importance of the permissive land use regulations in these cities, we estimate counterfactual output under the assumption that land use regulations in the South are set to the median city. This has the effect of making land use regulations in the southern cities more stringent and the housing supply more *inelastic*. Since local TFP in the southern cities grew rapidly over this period, this counterfactual housing supply elasticity results in higher housing prices and wages and lower employment in the South.

This has a large adverse effect on the aggregate growth rate. The growth rate of aggregate output would be almost 25 percent lower under this scenario. We get similar results when we set land regulations in the other large cities equal to that of the median US city.

We now turn to the case of imperfect mobility. As discussed above, we calibrate θ based on recent findings in Hornbeck and Moretti (2018), which estimate the long-run elasticity of local labor supply to be $1/\theta = 0.3$. Under this assumption, we find estimates of the effects of changes in land use regulations that are smaller relative to the perfect mobility case, as one might expect, but remain economically quite sizable.

²⁸ US GDP in 2009 was \$14.5 trillion so a GDP increase of 8.9 percent implies an additional aggregate income of \$1.95 trillion. Given a labor share of 0.65, this amounts to an increase of \$1.27 trillion in the wage bill, or \$8,775 additional salary per worker assuming a fixed number of workers. The salary increase would be smaller if more workers decide to enter the labor market in response to the higher salary or if there is immigration.

TABLE 6—REALLOCATION OF EMPLOYMENT: EFFECT OF MORE ELASTIC HOUSING SUPPLY IN NEW YORK, SAN FRANCISCO, AND SAN JOSE

	Percent change in rate of employment growth	
	Perfect mobility	Imperfect mobility
<i>Cities with largest increases</i>		
New York-Newark, NY-NJ-PA	1,010	318
San Francisco, CA	689	285
Flint, MI	479	161
San Jose, CA	243	108
Las Vegas, NV-AZ	-35	-12
<i>Cities with median change</i>		
Richmond-Petersburg, VA	-50	-17
<i>Cities with largest decreases</i>		
Binghamton, NY	-169	-57
Muncie, IN	-171	-57
Mansfield, OH	-214	-72
Youngstown-Warren, OH	-218	-73
Kokomo, IN	-295	-99

Notes: The sample includes 220 metropolitan areas observed in both 1964 and 2009. Entries represent the percent difference between the rate of employment growth under the counterfactual housing supply elasticity in New York, San Francisco, and San Jose and the observed rate of employment growth. Counterfactual housing supply elasticity for New York, San Francisco, and San Jose is the housing supply elasticity if New York, San Francisco, and San Jose adopted the housing regulations of the median US MSA. Column 1 assumes perfect mobility. Column 2 assumes imperfect mobility.

In particular, Table 5 shows that changing the housing supply regulation in New York, San Jose, and San Francisco to that in the median US city would increase the growth rate of aggregate output by 36.3 percent (second row). The net effect is that US GDP in 2009 would be 3.7 percent higher under this counterfactual, which translates into an additional \$3,685 in average wages for all workers, or an increase of \$0.53 trillion in the wage bill. The salary increase would be smaller if more workers decide to enter the labor market in response to the higher salary or if there is immigration.

Next, we turn to the implied effects on city size. Table 6 shows how the growth rate of employment would change with perfect mobility (column 1) or imperfect mobility (column 2). Focusing on the latter, by a vast margin, New York is the city that would experience the largest percentage increase in employment growth: a staggering 318 percent increase. San Jose and San Francisco growth would also accelerate significantly: 285 percent and 161 percent, respectively. Flint and Las Vegas would also benefit. The median city, Richmond, Virginia, would experience a slower growth (-17 percent). The bottom of Table 6 reports the cities that would experience the largest decline in growth in employment. This group mostly includes Rust Belt cities like Muncie, Indiana; Mansfield, Ohio; Youngstown, Ohio; and Kokomo, Indiana.

These changes in employment are economically very large. We stress that these are intended to be long-term benchmarks. They are based on the assumption that as the population expands in an area, local services also expand to keep the per capita availability of schools, parks, public transit, and other public amenities stable

at their current levels. Thus, one should not think of these counterfactuals as taking place overnight and holding fixed public services. Rather, one should think of these counterfactuals as taking place slowly over the long run, matched with a steady increase in the supply of public services so that the per capita level of public services is unchanged.

Welfare.—Welfare estimates are in columns 2 and 3 of Table 4 and Table 5. The second row in Table 4 shows that the effect of changing the housing supply regulation only in New York, San Jose, and San Francisco to that in the median US city would increase the growth rate of aggregate welfare by 51.8–55.8 percent under the assumption of perfect mobility. The corresponding estimate in the imperfect mobility case would be much smaller: 13.1 percent–14.1 percent.

In the case of southern cities, this counterfactual has a negative effect on the welfare growth. The effect on the growth rate of aggregate welfare is –44.7 percent to –48.6 percent under perfect mobility and –40.8 percent to –44.1 percent under imperfect mobility.

The effect on welfare is even larger than the effect on output because of the effect of higher housing prices in the southern cities on the average price of housing in all cities. We get similar results when we set land regulations in the other large cities equal to that of the median US city.

V. Extensions

In this section, we consider how our results change when we generalize the model.

A. Labor Demand Parameters

We remind the reader that the estimates of the contribution of specific cities to aggregate growth are based only on data on employment growth, and do not depend on the specific values we assume about the labor demand parameters. These parameters only matter when we isolate the contribution of amenities and housing prices, as well as when we measure the effect of the housing supply elasticity. We will check the sensitivity of these estimates to alternative assumptions about the labor demand parameters.

First, if we keep the assumption that the labor demand parameters are the same in all cities, the aggregate effect of changes in housing prices are not sensitive to the labor share, as long as the sum of the labor and the capital share is fixed. This is shown in Appendix Table A2. In Appendix Table A3, we alter the labor or capital share to vary their sum.²⁹

Second, we can drop the assumption that the labor demand parameters are the same in all cities. The labor demand parameters can differ if industries differ in

²⁹While results are not sensitive to changes in labor or capital share for a given degree of return to scale, they are quantitatively sensitive to the degree of decreasing return to scale. The closer the sum $\alpha + \eta$ is to 1, the larger the output gain. This makes intuitive sense, because the sum $\alpha + \eta$ governs the returns to scale. With $\alpha + \eta$ close to 1, our technology approaches constant returns to scale and there is the most productive cities attracting an increasingly larger share of the economic activity of the country.

terms of the importance of the resource (labor) affected by housing costs and amenities and the mix of industries differs across cities. Specifically, suppose that total output in a city is the sum of the output produced in different industries indexed by j , where the production function of industry j in city i is $Y_{ij} = A_{ij}L_{ij}^{\alpha_j}K_{ij}^{\eta_j}$. Note that the labor and capital shares are now indexed by industry, and the effective labor share of a city depends on the particular mix of industries in the city.

The local nominal wage is still given by $W_i = V \frac{P_i^\beta}{Z_i}$, but the equations for labor demand and aggregate output are different. Specifically, total employment in a city is given by $L_i \equiv \sum_j L_{ij} = \sum_j \left(\frac{\alpha^{1-\eta_j} \eta_j^{\eta_j}}{R^{\eta_j}} \cdot \frac{A_{ij}}{W_i^{1-\eta_j}} \right)^{\frac{1}{1-\alpha_i-\eta_j}}$ and aggregate output Y is implicitly defined by $1 = \sum_i \sum_j A_{ij}^{\frac{1}{1-\alpha_i-\eta_j}} \left(\frac{\alpha^{1-\eta_j} \eta_j^{\eta_j}}{\bar{\alpha} R^{\eta_j} Y} \cdot \frac{\bar{Q}}{Q_i} \right)^{\frac{1-\eta_j}{1-\alpha_i-\eta_j}}$, where $\bar{\alpha} \equiv \sum_i \sum_j \frac{Y_{ij}}{Y} \cdot \alpha_j$ is the aggregate labor share. Finally, aggregate utility is still given by $V = \bar{\alpha} \frac{Y}{\bar{Q}}$, but note that the relevant labor share is the *aggregate* labor share $\bar{\alpha}$.

This generalization has no effect on the estimates of the contribution of specific cities to aggregate growth, but it potentially changes the estimated effect of changes in housing prices. To gauge this, we use data on total employment and the labor share in each one-digit industry in each city in 1964 and 2009 to estimate the contribution of changes in housing prices on aggregate growth.³⁰ When we hold the distribution of housing prices fixed under this parameterization of the labor share, the growth rate of GDP is 105.6 percent higher than in the data. Remember that our baseline estimate where we assume the labor share is the same in all cities is that the growth rate of GDP is 103.5 percent higher if the distribution of housing prices is fixed, so the results are virtually identical.

B. Specialization by Cities

The baseline model assumes that output of a city is a perfect substitute for the products made by other cities. Suppose instead that each city makes a differentiated product with a production function given by $Y_i = A_i L_i$. The demand for the product of each city is determined by utility per capita defined as $\left(\sum_i Y_{ij}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{(1-\beta)\sigma}{\sigma-1}} \cdot \frac{h_j^\beta}{Z_j}$ where Y_{ij} denotes consumption of city i 's output in city j , h_j is per capita housing, and Z_j is the value of amenities in city j . Each city sells its product to other cities and buys the products made by the other cities.

We initially assume that shipping costs are zero. Indirect utility is then given by $V = W_i/Q_i$, where the local price is now given by $Q_i = Z_i/(C \cdot P_i^\beta)$, and $C \propto \left(\sum_j (A_j/W_j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ is the cost-minimizing price of a unit of the CES aggregate of the output good $\left(\sum_i Y_i^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$. Because shipping costs are zero, the price of

³⁰We use data on employment by one-digit industry and city from the County Business Pattern database. The data on the labor share for each one-digit industry is from Close and Shulenburg (1971) for 1964 and BEA (2013) for 2009.

the differentiated output goods are the same in all cities so C is also the same in all cities (and we normalize this price to one). All the other assumptions are the same.

Employment in a city is given by $L_i \propto \frac{A_i^{\sigma-1}}{V^{\sigma-1} Q_i^\sigma}$, aggregate output by $Y = \left(\sum_i A_i^{\sigma-1} (\bar{Q}/Q_i)^{\sigma-1} \right)^{\frac{1}{\sigma-1}}$, and aggregate utility by $V = \frac{\sigma-1}{\sigma} \cdot \frac{Y}{\bar{Q}}$. The equations for labor demand, aggregate output, and aggregate utility in this model with trade are the same as in our baseline model (without capital and where commercial land is normalized to one in all cities) where we substitute $\frac{1}{1-\alpha}$ with $\sigma-1$. Therefore, as long as we pick σ such that $\sigma-1 = \frac{1}{1-\alpha}$, all the results will be identical.

We now drop the assumption of costless shipment of goods across cities. Suppose τ_{ij} denotes the iceberg cost of shipping from city i to j . The price of the output goods now varies across cities because of the trade frictions. The local price is still given by $Q_i = Z_i / (C_i \cdot P_i^\beta)$, where the key difference is that C_i now varies across cities because of shipping costs. Holding amenities and housing prices fixed, C_i will be lower in more productive cities as long as shipping costs within a city are lower than shipping costs between cities. In addition, assuming that shipping costs are a function of physical distance between cities, C_i will also be lower in cities that are geographically proximate to high productivity cities.

Allowing for trade frictions has the following implications. First, the local nominal wage is still given by $W_i = \frac{V \cdot Z_i}{C_i \cdot P_i^\beta}$, but now differences in the nominal wage will reflect differences in the price of traded goods as well as differences in housing prices and amenities. Holding everything equal, this effect will show as *lower* wages in more productive cities and in cities close to productive cities.

Second, differences in prices of traded goods will also affect local employment. Local labor demand is now given by $L_i \propto \frac{A_i^{\sigma-1}}{V^{\sigma-1} Q_i^\sigma} \left(\sum_j \frac{Y_j}{Y} \tau_{ij}^\sigma C_j^\sigma \right)$. The effect of trade frictions shows up in two places. First, cities that are physically closer to more productive cities have a lower price index because imports from more productive cities will be cheaper. As discussed earlier, this lowers the local nominal wage, which also increases local employment. Second, trade barriers on a city's *exports* also affect local employment (this is the second term in the labor demand equation). Specifically, a city that faces larger trade barriers on its exports will have lower employment. These are the agglomeration forces modeled by Krugman (1991) and most recently by Allen and Arkolakis (2014).

Turning to aggregate output and welfare, aggregate output is given by $Y = \left(\sum_i A_i^{\sigma-1} \left(\frac{\tau_{ij}}{\bar{\tau}} \cdot \frac{\bar{Q}}{Q_i} \right)^{\sigma-1} \right)^{\frac{1}{\sigma-1}}$ where $\bar{\tau} \equiv \sum_j \sum_i \tau_{ji} \cdot \frac{C_{ji} Y_{ji}}{Y}$ is the revenue-weighted average of the iceberg trade cost.³¹ Aggregate utility is still given by $V = \frac{\sigma-1}{\sigma} \cdot \frac{\bar{\tau} Y}{\bar{Q}}$, where the only difference is that labor income is the product of the labor share and aggregate output net of the losses due to shipping costs.

³¹ We normalize the price of the CES bundle of the differentiated varieties in city i to 1.

The main empirical implication for the data inference exercise is that we can no longer back out the value of local amenities from data on nominal wages and housing prices. Put differently, the dispersion in the residual of the nominal wage (after controlling for housing prices) in Figure 5 reflects differences across cities in the price of the traded good as well as the dispersion in local amenities. In the previous section, we show that the increased dispersion in the nominal wage is entirely driven by the increased dispersion in housing prices (rather than dispersion in the residual of the nominal wage). Furthermore, our calculations focus on the effect of changes in housing prices and the effect of the housing supply elasticity on aggregate output and welfare, and these two calculations are the same in a model where cities trade subject to shipping costs.

C. Ownership of Housing Stock

We have assumed that workers do not own the housing stock so that an increase in average housing prices lowers welfare holding aggregate output fixed. Suppose we assume instead that the housing stock is owned by the workers in equal proportions, irrespective of where they live. For example, suppose that workers own equal shares of a mutual fund that owns all the housing in the United States. All the equations are the same, except that welfare is given by $V \propto (Y + \sum_i L_i h_i P_i) \cdot (\sum_i L_i \cdot P_i^\beta / Z_i)^{-1}$ where h_i denotes per capita housing consumption in city i . After imposing the condition that the share of nominal expenditures on housing is equal to β , the change in housing prices has the same effect on nominal income as on the average price of housing. In this case, changes in housing prices only affect welfare through the effect of the dispersion of the nominal wage on aggregate output, but changes in the average price of housing have no effect on welfare.

D. Endogenous TFP and Amenities

We have so far assumed that TFP and amenities are exogenous. However, a large literature in urban economics posits that city size (or density) also affects TFP due to agglomeration economies. Similarly, noise, traffic, and pollution can worsen when cities get larger. Urban amenities such as the variety of restaurants and cultural events can also change with city size. Allowing for these forces makes our estimates of the effect of changes in housing prices more complicated, as local TFP and amenities can change endogenously in response to changes in housing prices via their effect on city size.

For our purposes, the question is not whether TFP and amenities change endogenously when city size changes, but whether the elasticities of agglomeration and amenities to city size vary with the size of the city. If the elasticity of agglomeration to city size is the same in all cities, then the gains in TFP in cities that grow is exactly offset by TFP losses in cities that shrink. Similarly, if the elasticity of congestion or urban amenities to city size is the same in all cities, then the endogenous change in amenities in cities that increase is exactly offset by the change in amenities in cities that shrink. In both cases, our estimates of the aggregate implications of changes in housing prices are unchanged because the endogenous gain in TFP and amenities is exactly offset by the endogenous loss in TFP and amenities.

Empirically, the evidence from US manufacturing in Kline and Moretti (2014) and Albouy (2008) suggests that the elasticity of endogenous agglomeration to city size is the same in large versus small cities. As for urban amenities, Albouy (2008) shows that the quality of life in a city is positively correlated with the city population, but when natural amenities such as weather and coastal location are controlled for, there is no relationship between city population and the quality of life. This suggests that cities with better natural amenities are bigger (just as predicted by the equilibrium expression for city size in equation (6)), but endogenous amenities are not significantly better or worse in large cities compared to small cities.

If we believe the estimates in Kline and Moretti (2014) and Albouy (2008), the aggregate effect of changes in housing prices is the same in a model where amenities and TFP are endogenous as in a model where they are not. However, the estimates of the elasticity of TFP and amenities to city size are based on ranges of city size historically observed in the US data, and there is no guarantee that the same estimates extend to variation in city size that are significantly larger than the ones observed in the data.

VI. Policy Implications

Housing regulations have historically been perceived as having only local effects—affecting the quality of life and housing prices in the local community—and for this reason they have always been set at the municipality level. Our main point, however, is that local housing supply constraints can also have large effects on other cities. Specifically, we find that a major impediment to a more efficient allocation of labor across US cities is the constraint to housing supply in high TFP cities.

Although labor productivity and labor demand grew most rapidly in New York, San Francisco, and San Jose, thanks to the concentration of human capital intensive industries like high tech and finance, growth in these three cities had limited benefits for the United States as a whole. In the presence of strong labor demand, tight housing supply constraints effectively limited employment growth in these cities. We estimate that holding constant land availability, lowering regulatory constraints in New York, San Francisco, and San Jose cities to the level of the median city would increase aggregate output and welfare growth.

For example, Silicon Valley—the area between San Francisco and San Jose—has some of the most productive labor in the globe. But, as Glaeser (2014) puts it, “by global urban standards, the area is remarkably low density” due to land use restrictions. In a region with some of the most expensive real estate in the world, surface parking lots, one-story buildings and underutilized pieces of land are still remarkably common due to land use restrictions. While the region’s natural amenities—its hills, beaches and parks—are part of the attractiveness of the area, there is enough underutilized land within its urban core that housing units could be greatly expanded without any reduction in natural amenities.³² Our point is that a first-order effect of more housing in Silicon Valley is to raise income and welfare of *all* US workers.

³²Of course, supply of local public services—schools, police, public transit—would have to be expanded to keep the *per capita* availability constant.

In principle, one possible way to minimize the negative externality created by housing supply constraints in high TFP cities would be for the federal or state government to constrain US municipalities' ability to set land use regulations. If such policies have meaningful nationwide effects, then the adoption of federal or state standards intended to limit negative externalities may be in the aggregate interest. These types of standards exist in other countries. The state of California is currently debating a state statute that would significantly curtail municipalities' ability to deny or delay approval of housing projects that meet certain criteria (SB 35-2017).

An alternative is the development of public transportation that links local labor markets characterized by high productivity and high nominal wages to local labor markets characterized by low nominal wages. For example, a possible benefit of the high speed train currently under construction in California is to connect low wage cities in California's Central Valley—Sacramento, Stockton, Modesto, Fresno—to high productivity jobs in the San Francisco Bay Area. This could allow the labor supply to the San Francisco economy to increase overnight without changing San Francisco housing supply constraints. An extreme example is the London metropolitan area. A vast network of trains and buses allows residents of many cities in southern England—including far away cities like Reading, Brighton, and Bristol—to commute to high TFP employers located in downtown London. Another example is the Tokyo metropolitan area. While London and Tokyo wages are significantly above the UK and Japan averages, they would arguably be even higher in the absence of these rich transportation networks. Our argument suggests that UK and Japan GDP are significantly larger due to the transportation network.

APPENDIX

In this Appendix, we describe where each variable used in the paper comes from. We measure average wages in a county or in a country-industry cell by taking the ratio of total wage bill in private sector industries and total number of workers in private sector industries using data from the County Business Patterns (CBP) for 1964–1965 (referred to as 1964) and 2008–2009 (referred to as 2009). To increase sample size and reduce measurement error, we combine 1964 with 1965 and 2008 with 2009. 1964 is the earliest year for which CBP data are available at the county-industry level. Data on total employment by county are never suppressed in the CBP. By contrast, data by county and industry are suppressed in the CBP in cases where the county-industry cell is too small to protect confidentiality. In these cases, the CBP only provides a range for employment. We impute employment in these cases based on the midpoint of the range. We aggregate counties into MSAs using a crosswalk provided by the census based on the 2000 definition of MSA.

The main strength of the CBP is a fine geographical-industry detail and the fact that data are available for as far back as 1964.³³ But CBP is far from ideal. The main limitation of the CBP data is that it does not provide worker-level data on salaries, but only a county aggregate and therefore does not allow us to control for

³³Unfortunately, individual-level data on employment and salary with geocodes is not available from the Census of Population on a systematic basis until 1980. A third of metro areas are identified in the 1970 census.

changes in worker composition. We augment CBP data with information on worker characteristics from the Census of Population and the American Community Survey (ACS). Specifically, we merge 1964 CBP average wage by MSA to a vector of worker characteristics from the 1960 US Census of Population; we also merge 2009 CBP average wage by MSA to a vector of worker characteristics from the 2008 and 2009 ACS. These characteristics include: three indicators for educational attainment (high school dropout, high school, college), indicators for race, an indicator for gender, and age. We drop all cases where education is missing. In the small number of cases where one of the components of the vector other than education is missing, we impute it based on the relevant state average.

Because the census does not report information on union status, we augment our merged sample using information on union density by MSA from Hirsch, Macpherson, and Vroman (2001). Their data represent the percentage of the MSA nonagricultural wage and salary employees who are covered by a collective bargaining agreement. Their estimates for 1964 and 2009 are based on data from the Current Population Survey Outgoing Rotation Group (ORG) earnings files and the now discontinued BLS publication “Directory of National Unions and Employee Associations (Directory),” which contains information reported by labor unions to the federal government. The exact methodology is described in Hirsch, Macpherson, and Vroman (2001).³⁴

This allows us to estimate average residual wage in each MSA, defined as the average wage conditional on worker characteristics. Specifically, we estimate residual wages as $W_{ic} - X'_i b$, where W is average wage in the MSA, X is the vector of average worker characteristics in the MSA, and b is a vector of coefficients on worker characteristics from individual-level regressions estimated on a nationwide sample in 1964 and 2009 based on CPS data. The coefficients for 1964 are: high school or more: 0.44; college or more: 0.34; female: -1.13 ; nonwhite: -0.44 ; age: 0.004; union: 0.14. The coefficients for 2009 are: high school or more: 0.50; college or more: 0.51; female: -0.45 ; nonwhite: -0.07 ; age: 0.007; union: 0.14. Because a union identifier is not available in the 1964 CPS, the 1964 regression assumes that the coefficient on union is equal to the coefficient from 2009, which is estimated to be equal to 0.14.

For 2009, we can compare the wage residuals estimated in our approach with those that one would obtain from individual-level data. (Of course we can't do this for 1964, because we don't have micro data in that year). Figure A1 shows that while noisy, our imputed wage residuals do contain signal. The correlation between the two measures of the residual wage is 0.75.

We also estimated the residual wage as $W_{ic} - X'_i b_s$, where b_s is a vector of coefficients on worker characteristics from individual-level regressions, which is allowed to vary across states. The correlation in 2009 increases only marginally to 0.78.

Data on housing costs are measured as median annual rent from the 1960, 1970 US Census of Population and the 2008 and 2009 American Community Survey. For

³⁴For 1964, estimates are calculated based on figures in the BLS directories, scaled to a level consistent with CPS estimates using information on years in which the two sources overlap. Only state averages are estimated in 1964. Thus, in 1964, we assign union density to each MSA based on the state average.

TABLE A1—SUMMARY STATISTICS

	1964 average	2009 average
Average annual salary—private sector workers	25,538 (3,868)	29,018 (5,278)
Average annual rent	4,770 (932)	6,553 (1,826)
Private sector employment	144,178 (294,016)	377,071 (604,448)
Private sector wage bill (billion)	4.04 (8.95)	14.04 (25.5)
High school dropout	0.59 (0.11)	0.10 (0.05)
High school or more	0.40 (0.08)	0.90 (0.04)
College or more	0.07 (0.02)	0.26 (0.07)
Hispanic	0.03 (0.05)	0.10 (0.07)
Nonwhite	0.09 (0.11)	0.22 (0.15)
Age	28.1 (3.3)	39.9 (0.9)
Female	0.51 (0.01)	0.51 (0.01)
Union	0.26 (0.12)	0.11 (0.06)
Number of cities	220	220

Note: Mean and standard error (in parentheses) in 220 MSAs in each year.

TABLE A2—ESTIMATES BASED ON ALTERNATIVE PARAMETERS: OUTPUT

	Baseline				
	$\alpha = 65$	$\alpha = 6$	$\alpha = 55$	$\alpha = 70$	$\alpha = 75$
	$\eta = 25$	$\eta = 30$	$\eta = 35$	$\eta = 20$	$\eta = 15$
<i>Perfect labor mobility</i>					
Hold housing prices/amenities fixed	112.7	95.5	79.7	131.2	151.0
Hold housing prices fixed	103.5	87.1	72.2	121.3	140.6
Hold housing prices fixed, NY-SF-SJ	101.6	84.0	68.4	120.9	141.9
Set elasticity to median, NY-SF-SJ	86.1	74.6	63.6	98.1	110.7
	Baseline				
	$\alpha = 65$	$\alpha = 68$	$\alpha = 62$	$\alpha = 65$	$\alpha = 65$
	$\eta = 25$	$\eta = 25$	$\eta = 25$	$\eta = 28$	$\eta = 22$
<i>Imperfect labor mobility</i>					
Hold housing prices/amenities fixed	112.7	251.8	59.5	231.0	65.5
Hold housing prices fixed	103.5	280.6	55.3	219.6	58.9
Hold housing prices fixed, NY-SF-SJ	101.6	250.9	48.9	228.3	54.6
Set elasticity to median, NY-SF-SJ	86.1	171.0	48.8	158.8	55.3

Note: Column 1 is from Table 3 (column 1).

1964, we linearly interpolate census data between 1960 and 1970. Data for 2009 are from individual-level data from the American Community Survey. To get a more

TABLE A3—ESTIMATES BASED ON ALTERNATIVE PARAMETERS: WELFARE

	Baseline				
	$\alpha = 65$	$\alpha = 6$	$\alpha = 55$	$\alpha = 70$	$\alpha = 75$
	$\eta = 25$	$\eta = 30$	$\eta = 35$	$\eta = 20$	$\eta = 15$
<i>Perfect labor mobility</i>					
Hold housing prices/amenities fixed	115.0	100.0	86.8	131.0	147.9
Hold housing prices fixed	86.8	74.4	63.0	101.1	114.5
Hold housing prices fixed, NY-SF-SJ	73.7	62.8	53	85.9	99.4
Set elasticity to median, NY-SF-SJ	51.8	49.6	47.4	53.7	55.6
	Baseline				
	$\alpha = 65$	$\alpha = 68$	$\alpha = 62$	$\alpha = 65$	$\alpha = 65$
	$\eta = 25$	$\eta = 25$	$\eta = 25$	$\eta = 28$	$\eta = 22$
<i>Imperfect labor mobility</i>					
Hold housing prices/amenities fixed	115.0	233.0	67.6	215.4	73.2
Hold housing prices fixed	86.8	189.0	48.3	173.3	52.7
Hold housing prices fixed, NY-SF-SJ	73.7	172.6	40.7	156.8	44.3
Set elasticity to median, NY-SF-SJ	51.8	62.5	43.7	61.3	44.9

Note: Column 1 is from Table 3 (column 2).




FIGURE A1. ESTIMATED 2009 AVERAGE WAGE RESIDUAL VERSUS ACTUAL 2009 AVERAGE WAGE FROM INDIVIDUAL-LEVEL DATA

Notes: Each dot is an MSA. The x-axis reports average residuals by MSA from an individual-level regression based on individual-level data from the Census of Manufacturers. The y-axis has residuals based on CBP data used in the main analysis. The employment weighted correlation is 0.75.

precise estimate, we combine 2008 and 2009. Because rents may reflect a selected sample of housing units, we have re-estimated all our models using average housing prices. Results are essentially unchanged. Results are available upon request.

Our sample consists of 220 MSAs with non-missing values in 1964 and 2009. These cities account for 71.6 percent of US employment in 1964 and 72.8 percent in 2009. They account for 74.3 percent of US wage bill in 1964 and 76.3 percent

in 2009. The average city employment is 144,178 in 1964 and 377,071 in 2009. Table A1 presents summary statistics.

Data on housing supply elasticities, land use regulations, and land availability are from Saiz (2010). They are intended to measure variation in elasticity that arises both from political constraints and geographical constraints. In 19 cities, Saiz data are missing. In those cases, we impute elasticity based on the relevant state average.

REFERENCES

- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (6): 2127–89.
- Albouy, David. 2008. "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas." National Bureau of Economic Research (NBER) Working Paper 14472.
- Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *Quarterly Journal of Economics* 129 (3): 1085–1140.
- Barro, Robert J., and Xavier Sala-i-Martin. 1992. "Convergence." *Journal of Political Economy* 100 (2): 223–51.
- Basu, Susanto, and John G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy* 105 (2): 249–83.
- Baum-Snow, Nathaniel, and Ronni Pavan. 2012. "Understanding the City Size Wage Gap." *Review of Economic Studies* 79 (1): 88–127.
- Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2014. "Productive Cities: Sorting, Selection, and Agglomeration." *Journal of Political Economy* 122 (3): 507–53.
- Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin. 2011. "Finance and Development: A Tale of Two Sectors." *American Economic Review* 101 (5): 1964–2002.
- Bureau of Economic Analysis. 2013. National Economic Accounts. <https://www.bea.gov/data/economic-accounts/national>.
- Busso, Matias, Jesse Gregory, and Patrick Kline. 2013. "Assessing the Incidence and Efficiency of a Prominent Place Based Policy." *American Economic Review* 103 (2): 897–947.
- Close, Frank A., and David E. Shulenburger. 1971. "Labor's Share by Sector and Industry, 1948–1965." *ILR Review* 24 (4): 588–602.
- Davis, Morris A., and François Ortalo-Magné. 2011. "Household Expenditures, Wages, Rents." *Review of Economic Dynamics* 14 (2): 248–61.
- De La Roca, Jorge, and Diego Puga. 2017. "Learning by Working in Big Cities." *Review of Economic Studies* 84 (1): 106–42.
- Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000." *American Economic Review* 106 (3): 479–524.
- Diamond, Rebecca. 2017. "Housing Supply Elasticity and Rent Extraction by State and Local Governments." *American Economic Journal: Economic Policy* 9 (1): 74–111.
- Duranton, Gilles, Ejaz Ghani, Arti Grover Goswami, and William Kerr. 2015. "The Misallocation of Land and Other Factors of Production in India." World Bank Policy Research Working Paper 7221.
- Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny. 2014. "Spatial Sorting." *Journal of Political Economy* 122 (3): 554–620.
- Fajgelbaum, Pablo D., Eduardo Morales, Juan Carlos Suárez Serrato, and Owen M. Zidar. 2015. "State Taxes and Spatial Misallocation." National Bureau of Economic Research (NBER) Working Paper 21760.
- Ganong, Peter, and Daniel Shoag. 2017. "Why Has Regional Income Convergence in the US Declined?" *Journal of Urban Economics* 102: 76–90.
- Garicano, Luis, Claire Lelarge, and John Van Reenen. 2016. "Firm Size Distortions and the Productivity Distribution: Evidence from France." *American Economic Review* 106 (11): 3439–79.
- Gennaioli, Nicola, Rafael La Porta, Florencio Lopez De Silanes, and Andrei Shleifer. 2014. "Growth in Regions." *Journal of Economic Growth* 19 (3): 259–309.
- Glaeser, Edward. 2014. "Land Use Restrictions and Other Barriers to Growth." *CATO Online Forum*, December 1. <https://www.cato.org/publications/cato-online-forum/land-use-restrictions-other-barriers-growth>.

- Glaeser, Edward L., and Joseph Gyourko. 2003. "The Impact of Building Restrictions on Housing Affordability." *Federal Reserve Bank of New York Economic Policy Review* 9 (2): 21–39.
- Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy* 113 (2): 345–75.
- Glaeser, Edward L., Joseph Gyourko, and Raven E. Saks. 2006. "Urban Growth and Housing Supply." *Journal of Economic Geography* 6 (1): 71–89.
- Glaeser, Edward L., and David C. Maré. 2001. "Cities and Skills." *Journal of Labor Economics* 19 (2): 316–42.
- Glaeser, Edward L., and Bryce A. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65 (3): 265–78.
- Gourio, François, and Nicolas Roys. 2014. "Size-Dependent Regulations, Firm Size Distribution, and Reallocation." *Quantitative Economics* 5 (2): 377–416.
- Greenwood, Jeremy, Juan M. Sanchez, and Cheng Wang. 2013. "Quantifying the Impact of Financial Development on Economic Development." *Review of Economic Dynamics* 16 (1): 194–215.
- Guner, Nezih, Gustavo Ventura, and Yi Xu. 2008. "Macroeconomic Implications of Size-Dependent Policies." *Review of Economic Dynamics* 11 (4): 721–44.
- Henderson, J. Vernon. 1982. "Systems of Cities in Closed and Open Economies." *Regional Science and Urban Economics* 12 (3): 325–50.
- Henderson, J. Vernon, and Yannis M. Ioannides. 1981. "Aspects of Growth in a System of Cities." *Journal of Urban Economics* 10 (1): 117–39.
- Hirsch, Barry T., David A. Macpherson, and Wayne G. Vroman. 2001. "Estimates of Union Density by State." *Monthly Labor Review* 86 (7): 51–55.
- Hopenhayn, Hugo, and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy* 101 (5): 915–38.
- Hornbeck, Richard, and Enrico Moretti. 2018. "Who Benefits from Productivity Growth? Direct and Indirect Effects of Local TFP Growth on Wages, Rents, and Inequality." http://faculty.chicagobooth.edu/richard.hornbeck/research/papers/TFP_HM.pdf.
- Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4): 1403–48.
- Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation: Dataset." *American Economic Journal: Macroeconomics*. <https://doi.org/10.1257/mac.20170388>.
- Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics* 129 (1): 61–103.
- Kline, Patrick. 2010. "Place Based Policies, Heterogeneity, and Agglomeration." *American Economic Review* 100 (2): 383–87.
- Kline, Patrick, and Enrico Moretti. 2014. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics* 6 (1): 629–62.
- Krugman, Paul. 1991. "Increasing Returns and Economic Geography." *Journal of Political Economy* 99 (3): 483–99.
- Lewbel, Arthur, and Krishna Pendakur. 2009. "Tricks with Hicks: The EASI Demand System." *American Economic Review* 99 (3): 827–63.
- Mayer, Christopher J., and C. Turiel Somerville. 2000. "Land Use Regulation and New Construction." *Regional Science and Urban Economics* 30 (6): 639–62.
- Midrigan, Virgiliu, and Daniel Yi Xu. 2014. "Finance and Misallocation: Evidence from Plant-Level Data." *American Economic Review* 104 (2): 422–58.
- Moll, Benjamin. 2014. "Productivity Losses from Financial Frictions: Can Self-Financing undo Capital Misallocation?" *American Economic Review* 104 (10): 3186–3221.
- Moretti, Enrico. 2012. *The New Geography of Jobs*. Boston: Houghton Mifflin Harcourt.
- Peters, Michael. 2016. "Heterogeneous Markups and Endogenous Misallocation." <https://mipeters.weebly.com/uploads/1/4/6/5/14651240/markupsmisallocation.pdf>.
- Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.
- Quigley, John M., and Steven Raphael. 2004. "Is Housing Unaffordable? Why Isn't It More Affordable?" *Journal of Economic Perspectives* 18 (1): 191–214.
- Redding, Stephen J., and Matthew A. Turner. 2014. "Transportation Costs and the Spatial Organization of Economic Activity." National Bureau of Economic Research (NBER) Working Paper 20235.
- Restuccia, Diego, and Richard Rogerson. 2013. "Misallocation and Productivity." *Review of Economic Dynamics* 16 (1): 1–10.
- Restuccia, Diego, and Richard Rogerson. 2017. "The Causes and Costs of Misallocation." *Journal of Economic Perspectives* 31 (3): 151–74.

- Restuccia, Diego, and Ral Santaaulàlia-Llopis.** 2015. "Land Misallocation and Productivity." https://www.economics.utoronto.ca/diegor/research/RS_paper.pdf.
- Saiz, Albert.** 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.
- Saks, Raven E.** 2008. "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth." *Journal of Urban Economics* 64 (1): 178–95.
- Serrato, Juan Carlos Suárez, and Owen Zidar.** 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review* 106 (9): 2582–2624.