Who Benefits From Productivity Growth? The Local and Aggregate Impacts of Local TFP Shocks on Wages, Rents, and Inequality *

Richard Hornbeck (Harvard and NBER)
Enrico Moretti (Berkeley and NBER)

May 2015 (Preliminary Draft)

Abstract

We estimate who benefits from local productivity growth. We begin by using confidential plant-level data to calculate changes in manufacturing productivity by United States metropolitan areas (MSAs), and then predict changes in MSA productivity based on MSAs' initial industry shares and subsequent national changes in productivity by industry. We find that local productivity growth benefits local landowners more than local workers, in percentage terms. Workers do benefit directly from local productivity growth, however, as there are generally positive net effects on real earnings. These local effects of TFP are different for skilled and unskilled workers: consistent with lower geographic mobility among less-educated workers, we estimate that TFP growth generates greater increases in the local number of more-educated workers and larger local wage increases for less-educated workers. Thus, local increases in TFP compress inequality at the local level (and local declines in TFP magnify inequality).

Geographic mobility induces general equilibrium effects from local changes in TFP, however, and so we then turn to the aggregate impacts of local changes in TFP. We find that a substantial portion (almost half) of the aggregate wage impacts accrue in cities indirectly affected through out-migration, particularly among more-mobile high-skill workers. By contrast, there is little aggregate impact on housing costs, as increases in cities directly impacted by TFP gains are mostly offset by losses in other cities. Thus, the aggregate economic incidence of local productivity shocks falls entirely on workers. Overall, the aggregate economic incidence of local productivity growth differs importantly from the local incidence, and in a manner more skewed toward more-mobile high-skill workers.

*For helpful comments and suggestions, we thank Chang-Tai Hsieh, Patrick Kline, and Martin Rotemberg. Joseph Root provided excellent research assistance, in addition to assistance from Georgios Angelis, Melissa Eccleston, Alonso Sanchez, and Alex Weckenman.

Economists have long considered productivity growth as a key determinant of economic progress and improvements in consumption levels. Indeed, in the long-run, increases in workers' living standards are naturally linked to productivity growth. Given the central role that productivity growth plays in economic development, it is not surprising that understanding its causes and effects has been central to modern economics.

When thinking about how productivity might affect worker welfare, however, there are at least two limitations with the existing literature. First, Total Factor Productivity (TFP) is difficult to quantify empirically and is often imputed using aggregate data. Second, and more fundamentally, most research on the effects of productivity growth is at the country level. This is probably not surprising, as there are huge and well-documented differences in TFP levels across countries and over time (see, e.g., Hall and Jones, 1999; Caseli, 1999). Thinking about productivity as geographically uniform within a country is unrealistic, however, as there are tremendous productivity differences within countries in the cross-section and in changes over time.[1] These differences reflect geographical heterogeneity across cities and regions in industry mix, local transportation infrastructure, R&D investment within industries, local economic policies, human capital spillovers, productivity spillovers, and many other local factors that influence local productivity beyond national trends.

Geographical variation in productivity levels and productivity shocks has important implications for the impacts of productivity on worker welfare. Localized productivity shocks might induce migration and, therefore, changes in housing costs (or other local goods) that profoundly alter the welfare consequences of productivity growth. For example, while New York and San Francisco have experienced rapid productivity growth over the past three decades, the cost of living has also increased substantially and has shifted some of the benefits of productivity growth from labor to incumbent landowners. While much recent attention has focused on changes over time in the share of output that accrues to labor and capital,[2] we know less about the incidence of productivity shocks once local prices are taken into account. Arguably, this has the potential to be much more consequential for the distribution of benefits generated by economic growth.[3] In addition, geographical mobility is different across skill groups, with high-skilled workers often more likely to move in response

[1] In the United States, for example, we estimate that manufacturing establishments in cities with high average productivity have TFP levels that are more than 50% higher than manufacturing establishments with the same capital stock and same employment in cities with low average productivity. Between 1980 and 1990, TFP increased by 18% in the average US city. This average increase masks profoundly different experiences across cities, with some cities enjoying TFP gains of 50% or more, and other cities experiencing stagnant or even declining TFP.

[2] See, for example, work by Blanchard (1997); Blanchard and Giavazzi (2003); Bentolila and Saint-Paul (2003); Jones (2005); Rodriguez and Jayadev (2010); Piketty (2014).

[3] A parallel question, which we do not address in this paper, is how much productivity increases benefit consumers as compared to capital owners.

to labor demand shocks than low skilled-workers. Thus, even skill-neutral shocks to local productivity can have differential incidence on different types of workers and, thus, impacts on inequality. These heterogeneous effects can be lost in more aggregate analysis.

In this paper, we estimate the incidence of localized productivity shocks in the United States, and how this incidence varies across skill groups. We focus on geographical differences in TFP changes across metropolitan areas, and investigate how much of the benefits from productivity increases accrue to workers (in the form of higher wages) and how much of the benefits accrue to incumbent landowners (in the form of higher housing rents). Unlike most of the previous literature, we *directly* measure local productivity changes using detailed plant-level data on the universe of manufacturing establishments in the United States.

We begin by laying out a simple theoretical framework that clarifies how in spatial equilibrium a localized TFP change might affect wages, cost of housing, and employment of different types of workers. An increase in TFP in a city results in an increase in labor demand in that city, and this triggers in-migration from other cities. As the price of land and other non-tradables adjusts, some of the benefits of the TFP increase accrue to landowners.[4] The effect on real wages, i.e., wages adjusted for local cost of living, can be quite different from the effect on nominal wages.

In addition, the ultimate impact on the equilibrium wage can be quite different across skill groups as a function of their geographical mobility. In particular, high-skilled individuals are generally more geographically mobile than low-skilled individuals (Wozniak, 2010; Malamud and Wozniak, 2012). Thus, the nominal wage of college graduates is expected to be less sensitive to local productivity shocks and their employment is expected to be more sensitive. This implies that a skill-neutral TFP shock can have profoundly different effects on equilibrium real wages and employment levels of skilled and unskilled workers, and therefore can affect income inequality at the local level.

Further, the model clarifies how the aggregate incidence of productivity shocks will diverge from the local incidence due to general equilibrium effects. In particular, through the geographic mobility of workers, some of the benefits of a TFP increase in one city will spread to other cities. For the same reason, the local impacts on inequality will diverge from the aggregate impacts on inequality due to general equilibrium effects of different magnitudes by worker type. The model provides a framework for calculating the aggregate incidence of local productivity shocks, drawing on estimates of the local incidence and the elasticity of labor demand.[5]

We use confidential plant-level data from the Census of Manufacturers to estimate Total

[4] To the extent that some workers are renters, gains to landowners need not pass on to workers.

[5] This approach complements recent work by Hsieh and Moretti (2015).

Factor Productivity (TFP) for each establishment in the US in 1980 and 1990. We aggregate TFP data to the level of metropolitan area (MSA), and link 1980-1990 TFP changes to MSA-level data on wages, housing costs, and employment between 1980 and 1990 (in our "short run" models) and between 1980 and 2000 (in our "long run" models). We focus on empirical specifications that instrument for changes in productivity, using the predicted change in MSA-level TFP based on each MSA's initial industry concentration and industry-level changes in productivity across all other MSAs (i.e., a "share-shift" instrument).

Our estimates indicate that local productivity growth benefits local landowners more than local workers, in percentage terms. Specifically, we find that a 1 percentage point increase in TFP between 1980 and 1990 in a city is associated with a 0.91% increase in annual nominal earnings per worker in that city in the same period, and a 1.86% increase in housing rents. Long-run effects exceed short-run effects, presumably because it takes time for workers to relocate in the new equilibrium and because of any agglomeration economies: a 1 percentage point increase in TFP between 1980 and 1990 in a city is associated with 1.45% increase in annual earnings between 1980 and 2000, and a 2.30% increase in rents.

Workers do benefit directly from local productivity growth, however, as the net effects on real earnings are generally positive. In particular, while housing rents increase by approximately twice the increase in annual earnings, workers spend less than half of their earnings on housing. Real wages increase by 0.76% in the long-run, assuming than 30% of earnings is spent on local goods whose prices move along with the rental cost of housing. Further, when workers own their house, they benefit from TFP growth through both the housing market and the labor market.[6]

Importantly, we estimate that the local effects of TFP shocks are different for skilled and unskilled workers. Consistent with lower geographic mobility among less-educated workers, we estimate that a 1 percentage point increase in TFP is associated with a 5.82% long-run increase in the number of college graduates in a city, and a 3.32% increase in the number of high school graduates. These estimates imply that the local supply of high school graduates is less elastic to changes in TFP than that of college graduates. Following a shift in labor demand, due to increased TFP, a less elastic labor supply should have larger wage effects. Indeed, we estimate that local increases in TFP result in larger wage increases for high school graduates (1.23%) than wage increases for college graduates (0.87%). As a consequence, local increases in TFP compress inequality at the local level, whereas local declines in TFP magnify inequality.

We also consider the reallocation of workers across sectors, within cities, and calculate the

[6]In particular, this is the case when workers own the houses before the TFP shock becomes known by the market.

implied short-run and long-run multiplier effects. Increases in manufacturing TFP directly impact the manufacturing sector, but also indirectly impact local non-manufacturing sectors by increasing demand for locally non-traded goods and services (Moretti, 2010). Indeed, policy efforts to support the growth of local manufacturing — or prevent the collapse of local manufacturing — are often justified by policymakers on these grounds and depend on the magnitude of the "local multiplier."[7] We estimate that for an increase in manufacturing TFP that creates one manufacturing job, there is an implied short-run increase of 1.24 non-manufacturing jobs and a long-run increase of 2.96 non-manufacturing jobs. This short-run estimate of the local multiplier is moderately smaller than previous estimates (Moretti, 2010), and the long-run estimate provides new evidence on how the local multiplier increases over time.

In the final part of the empirical analysis, we turn to calculating the aggregate impacts of local changes in TFP. While our initial analysis of micro-data provides estimates of local incidence, the theoretical framework clarifies how geographic mobility can affect the aggregate incidence of TFP shocks. Geographic mobility reduces the local economic gains from local increases in TFP, as worker inflows depress wages, but higher geographic mobility generates greater spillover benefits to other cities and raises the aggregate economic gains for those worker types. Overall, a large fraction of the effects of localized TFP shocks are not localized, as a substantial portion of the aggregate impacts from TFP increases occurs in cities other than the cities directly hit by changes in TFP from 1980 to 1990. Indeed, almost one-half of the overall increase in worker nominal earnings occurs in cities not directly impacted by TFP changes. There is little aggregate impact on housing costs, as the increases in cities directly impacted by TFP gains are mostly offset by losses in other cities.

The aggregate economic importance of general equilibrium effects is different, however, for high-skill and low-skill workers. For workers with a college degree, 59% of the overall increase in worker earnings occurs in cities not directly impacted, whereas the corresponding figure for high school graduates is 42%. This difference reflects the different propensity to move among these two groups. As a consequence, the aggregate incidence of TFP shocks differs importantly from the local incidence, and in a manner more skewed toward high-skill workers.

An important policy question is whether a city can affect the incidence of TFP shocks through local land-use regulation. Saiz (2010) and Glaeser and Gyourko (2005) have pointed out that cities differ widely in their elasticity of housing supply, both for geographical reasons

[7] For example, the state of Nevada recently provided the electric car manufacturer Tesla Motors an incentive package worth $400 million to attract Tesla's latest battery factory. A significant component of the state's justification for the incentive package was the multiplier effect. These types of incentives are common (Kline and Moretti, forthcoming).

and policy reasons, and have shown that housing prices are lower in cities where land-use regulations are more permissive and the housing supply is more elastic. Similarly, we find that a TFP shock has a larger effect on rents in cities with a more inelastic housing supply. We also find that nominal wages adjust almost proportionally, however, which is consistent with a high degree of labor mobility. Thus, on net, the local impacts of a TFP shock on real wages are similar across cities with different elasticities of housing supply. City governments therefore have limited policy influence over the fraction of economic benefits from productivity growth that go to local workers relative to local landowners. Further, from an aggregate perspective, restrictions on housing supply and worker mobility directly lower the aggregate economic gains from local increases in productivity.

The remainder of this paper is organized as follows. In Section I we present a simple theoretical framework. In Sections II and III we discuss the data and econometric models. The main results are in Sections IV and V. Section VI concludes.

I Theoretical Framework

In our empirical analysis, we estimate how an increase in city Total Factor Productivity (TFP) impacts nominal wages, housing costs, and real wages in that city as well as in other cities. This section presents a simple spatial equilibrium model of the labor market and housing market that is useful in guiding our empirical approach and interpreting our findings.

The model adopts the standard assumptions of most Rosen-Roback spatial equilibrium models. The setting and the specific functional form assumptions are similar to those in Moretti (2011).

We assume that each city is a competitive economy, producing a single output good Y that is traded on the international market at a given price, assumed to be 1. The production function for firm i located in city c is:

$$Y_{ic} = X_c N_{ic}^{(1-h)}, \tag{1}$$

where X_c is a city-specific productivity shifter and N_{ic} is labor input. Labor demand is derived from the usual first order conditions, and has the usual form:

$$w_c = x_c + log(1-h) - hn_c, \tag{2}$$

where w_c is the log nominal wage in city c, x_c is log productivity, and n_c is the log labor force.

Consider the case of two cities and a fixed number of workers. Utility of worker in city c depends on consumption of the traded good Y, housing H, the city's local amenities A, and an idiosyncratic location preference: $U_{ic} = Y_{ic}^{(1-\beta)} H_{ic}^{\beta} A_c u_{ic}$. In particular, we assume that worker i's relative preference for city a over city b $(u_{ia} - u_{ib})$ has a type 1 extreme value distribution with mean 0 and variance s^2. The parameter s governs the importance of idiosyncratic location preferences and, therefore, the degree to which workers are willing to change cities to arbitrage away differences in real wages. For smaller values of s, more workers are willing to move for a higher real wage. Alternatively, one can think of u as a moving cost.

The (log) indirect utility of worker i in city c is then:

$$v_{ic} = w_c - \beta r_c + a_c + e_{ic} \tag{3}$$

where r_c is the log of cost of housing, a_c is the log value of amenities, and $e_{ic} = \ln u_{ic}$. The parameter β measures the importance of housing consumption in utility, and equals the budget share spent on housing. Thus, the effect on indirect utility from a 1% increase in wages is larger than the effect from a 1% decrease in cost of housing.

We assume that workers and firms are mobile, locating wherever utility and profits are maximized. Worker i chooses city b, rather than city a, if and only if the strength of location preferences exceeds any real wage premium and higher amenity value: $e_{ib} - e_{ia} > (w_a - \beta r_a) - (w_b - \beta r_b) + (a_a - a_b)$. This implies that local labor supply in city b is:

$$w_b = -s + w_a + \beta(r_b - r_a) + (a_a - a_b) + s(n_b - n). \tag{4}$$

Intuitively, when location preferences are stronger (i.e., when s is larger), workers' location decisions are less sensitive to differences in real wages and the labor supply curve is steeper. When location preferences are weaker, workers' location decisions are more sensitive to differences in real wages and cities face a less steep labor supply curve. In the extreme case, when s is zero and there are no idiosyncratic location preferences, labor supply is infinitely elastic to any difference in real wages and the labor supply curve in equation 4 is horizontal.

We assume that each worker consumes one homogeneous housing unit and that the log price of housing is governed by:

$$r_c = k_c n_c. \tag{5}$$

This is simply the reduced-form relationship between the log cost of housing and the log number of residents in city c. The parameter k_c reflects differences in the elasticity of housing supply and varies across cities, reflecting differences in geographical constraints and local regulations on land development. In cities where the geography and regulatory structure make it relatively easy to build new housing, k_c is relatively smaller. In the extreme case where there are no constraints to building housing, the supply curve is horizontal and k_c is zero. In the extreme case where it is impossible to build new housing, the supply curve is vertical and k_c is infinite.[8]

I.A Local Effect of a Local Productivity Shock

We are interested in quantifying how a localized productivity shock in one city affects equilibrium wages, housing rents, and employment in that city and elsewhere.

We assume that the two cities are initially identical and that city b experiences an unexpected factor-neutral productivity shock that raises local Total Factor Productivity (TFP) in that city by an amount Δ. Using the model notation, if x_{b1} is TFP in city b before the shock, TFP after the shock is $x_{b2} = x_{b1} + \Delta$. Productivity in city a does not change.

In practice this productivity shock could be due to many reasons, including: a change in local transportation infrastructure, a change in local policies (e.g., the election of a pro-business mayor), or a localized technological shock. For a technological shock to be localized it would have to be either firm-specific (e.g., Google invents a new technology that raises productivity at Google and, therefore, average productivity where Google is located) or industry-specific. Since industries are not uniformly distributed over space, an industry-specific productivity shock will almost always have different impacts on cities that depend on the cities' industry mix. For example, if productivity increases in the IT industry due to a new technology, the effect is larger in San Jose where IT firms are concentrated than in cities where IT firms have little presence.

The increase in productivity in city b shifts local labor demand curve to the right, resulting in higher nominal wages and more employment. Higher employment leads to higher housing costs. Although no shock has hit city a directly, its labor market and housing market are indirectly affected through labor mobility. As employment declines in city a, the cost of

[8]For simplicity, we are ignoring the asymmetry between positive and negative shocks uncovered by Glaeser and Gyourko (2005). In our setting, this is unlikely to be a major problem, since most of the TFP shocks are positive (although they vary in relative magnitude across cities).

housing declines and real wages increase.

More precisely, in city b nominal wage and employment increase, respectively, by:

$$
\begin{aligned}
w_{b2} - w_{b1} &= \frac{(\beta(k_a + k_b) + h + s)}{\beta(k_a + k_b) + 2h + s}\Delta \\
&= \pi_2 \Delta,
\end{aligned} \tag{6}
$$

and

$$
\begin{aligned}
n_{b2} - n_{b1} &= \frac{\Delta}{\beta(k_a + k_b) + 2h + s} \\
&= \pi_1 \Delta.
\end{aligned} \tag{7}
$$

Because of in-migration, the cost of housing in city b also increases:

$$
\begin{aligned}
r_{b2} - r_{b1} &= \frac{\beta k_b}{\beta(k_a + k_b) + 2h + s}\Delta \\
&= \pi_3 \Delta.
\end{aligned} \tag{8}
$$

The magnitudes of these effects depend on the elasticities of labor supply and housing supply. From equation 7, the number of movers is larger when the elasticity of labor supply is high (i.e., s is small) and the elasticity of housing supply in b is high (i.e., k_b is small). Intuitively, a smaller s implies that idiosyncratic location preferences are less important, so workers are more mobile in response to real wage differentials. A smaller k_b means that city b can add more new housing units to accommodate in-migrants' demand with less of an increase in housing cost. For the same reasons, the smaller the elasticities of labor supply and housing supply in city b, the larger the increases in wages and cost of housing in city b (see equations 6 and 8). In the extreme case where the elasticity of housing supply is infinite ($k_b = 0$), housing costs do not change in city b.

The change in real wages in city b is obtained by comparing the increase in nominal wage with the increase in housing cost:

$$
\begin{aligned}
(w_{b2} - w_{b1}) - \beta(r_{b2} - r_{b1}) &= \frac{(\beta k_a + h + s)}{\beta(k_a + k_b) + 2h + s}\Delta \\
&= \pi_4 \Delta.
\end{aligned} \tag{9}
$$

The benefits from productivity growth are split between workers and landowners. Equation 9 clarifies that the incidence of the shock depends on which of the two factors — labor or land — is supplied more elastically at the local level.

For a given elasticity of housing supply, a lower local elasticity of labor supply implies that a larger fraction of the productivity shock in city b accrues to workers in city b, and a smaller fraction accrues to landowners in city b. When workers are *less* mobile, they capture more of the economic rent generated by the productivity shock (or are harmed more from a local decline in productivity). There will also be a smaller increase in real wages in the non-affected city (city a) when workers are less mobile, since worker movement is the channel that generates benefits for the non-affected city. In the extreme, if labor is completely immobile ($s = \infty$), then equation 9 becomes: $(w_{b2} - w_{b1}) - \beta(r_{b2} - r_{b1}) = \Delta$. In this extreme case, real wages in city b increase by the full amount of the productivity shock. Indeed, the benefit of the shock accrues entirely to workers in city b. Intuitively, when labor is a fixed factor, workers in the city directly impacted by the shock will capture the full economic rent generated by the shock.

For a given elasticity of labor supply, a lower elasticity of housing supply in city b relative to city a (k_b bigger than k_a) implies that a larger fraction of the productivity shock in city b accrues to landowners in city b and a smaller fraction accrues to workers. When housing supply is more inelastic, housing quantity adjusts less in city b following the productivity shock and housing prices increase more. In the extreme case where housing supply in city b is fixed ($k_b = \infty$), the entire productivity increase is capitalized into land values in city b.[9]

In summary, the model clarifies how a TFP shock affects local nominal wages, rents, real wages, and employment. Our empirical analysis will quantify the effect of a TFP shock on nominal wages, employment, housing rents, and real wages. We will also test these comparative statics. In particular, we will test whether the effect of a TFP shock varies as a function of labor supply and housing supply elasticities, as predicted by the model, using variation across cities in the elasticity of housing supply and variation across skill groups in geographical mobility.

I.B Aggregate Effect of a Local Productivity Shock

In our setting, labor markets are partially integrated across geographic space. This implies that the effects of a localized TFP shock in city b are not limited to that particular city, but spread to other cities through general equilibrium effects. Although city a does not

[9]We have begun by assuming that TFP is exogenous. In practice, it is plausible that TFP is endogenous to some extent. For example, a large literature in urban and regional economics posits that productivity is casually increasing in city density (e.g., $X_c = f(N_c)$). In this case, impacts of a productivity shock on city size would induce an additional feedback effect through further increases in productivity. In this case, our estimates of equations 6, 7, 8, and 9 should be interpreted as the effect of the original TFP shock plus any additional endogenous increase in TFP.

experience any *direct* shock to its economy, it is *indirectly* affected by the shock that hits city b. In our model, the mechanism through which shocks spread is labor mobility: the increase in city b TFP causes workers to leave city a for city b. In turn, for a given labor demand elasticity and housing supply elasticity, there are changes in the nominal wages and rents in city a.

In particular, the decrease in employment in city a is, by construction, equal to the increase in city b:

$$n_{a2} - n_{a1} = -(n_{b2} - n_{b1}), \tag{10}$$

and it is determined by equation 7. This result follows from the assumption of a fixed number of workers in the economy and, for notional simplicity, the assumption that city a and city b are initially of the same size. As workers move from city a to city b, nominal wages necessarily increase in city a and housing costs decrease. The net effect is an increase in real wages in city a:

$$(w_{a2} - w_{a1}) - \beta(r_{a2} - r_{a1}) = \frac{(\beta k_a + h)}{\beta(k_a + k_b) + 2h + s}\Delta. \tag{11}$$

The increase in real wages in city a is not necessarily the same as the increase in city b. Comparing equations 9 and 11, the effect of a productivity shock is generally larger in the city directly hit by the shock, as one would expect in a model with less than perfect mobility. Indeed, real wages will only be completely equalized if there is perfect labor mobility, i.e., in the absence of location preferences ($s = 0$). Intuitively, city b is directly affected by the productivity shock and real wages are only partially bid up in city a because some workers have a location preference to remain in city a. Only the marginal worker is indifferent between the two cities, in equilibrium, and only with no idiosyncratic location preferences would *all* workers be indifferent between the two cities.

In the simplified case of two cities, the indirect effects on city a are concentrated and large. In reality, however, there are many possible cities of origin for workers who move to city b. Therefore, in reality, the indirect effects on each of these cities are diffused and small. While the indirect effects are spread amongst many cities, their sum is still potentially large. The magnitude of the indirect effects on wages and rents in each origin city are not the same for all cities, as they depend on local elasticities of housing supply.[10]

[10]In principle, there could also be general equilibrium effects across the entire world, although labor markets are much more integrated within the United States than across international borders. Quantifying these general equilibrium spillovers across the entire world is beyond the scope of our analysis.

In a later section of the empirical analysis, we will quantify the indirect effects of TFP shocks. For each city experiencing a TFP shock, we estimate the change in employment. Under some assumption on mobility patterns, then calculate the indirect effect on employment levels in all other cities.[11] Given parameter values for the elasticity of labor demand and city-specific elasticities of housing supply, we can then calculate the corresponding effects on wages and rents in these other cities.

In related work, Hsieh and Moretti (2015) use an alternative approach to investigate the aggregate effects of localized shocks. There are three main differences in our approaches. First, they estimate TFP shocks from the first order conditions of the model, whereas we directly estimate TFP shocks from surveyed output and input data from individual establishments. Second, they estimate the effects of local shocks by calibrating the model, whereas we estimate empirically how local shocks affect local wages and rents. Third, they focus on the aggregate effects on output, whereas we focus on wages and rents that determine the incidence of productivity shocks.

II Data

Our dataset combines measures of local Total Factor Productivity (TFP) and local labor market outcomes.

II.A Total Factor Productivity

We measure manufacturing productivity using plant-level data from the Census of Manufacturers (CMF). The CMF contains plant-level data at five-year intervals on all manufacturing plants' employment, capital stocks, materials, and total value of shipments. We estimate TFP in each MSA and year, following the same approach as that used by Greenstone, Hornbeck and Moretti (2010). We assume that each plant p in year t uses the following Cobb-Douglas technology:

$$Y_{pt} = A_{pt} L_{pt}^{\beta_1} K_{pt}^{B\beta_2} K_{pt}^{E\beta_3} M_{pt}^{\beta_4}, \tag{12}$$

where Y is the total value of shipments minus changes in inventories, A is Total Factor Productivity (TFP), L is total labor hours, K^B is total book value of building capital stock, K^M is total book value of machinery and equipment capital stock, and M is value of material

[11]We will make these calculations under two alternative assumption on mobility flows across cities. As a benchmark, we will assume that workers moving to a particular city are drawn from all other cities in proportion to their number of workers. As an alternative, we will assume that workers moving to a particular city are drawn from all other cities in proportion to actual historical city-to-city flows.

inputs.[12] The regressions are weighted by plant output.[13]

Using plant-level data for each year separately, we regress log output on log materials, log labor, log capital, and MSA fixed effects. The estimated MSA fixed effects reflect average Total Factor Productivity (TFP) in each MSA and year. Plant-level data offer advantages in measuring local productivity differences, though there are known challenges in estimating Total Factor Productivity (TFP). The estimated coefficients on each input may be biased due to plants' endogenous input choices. For example, unobserved demand shocks are likely to affect input utilization, and this raises the possibility that the estimated coefficients on inputs are inconsistent (see, e.g., Griliches and Mairesse 1995). This has been a topic of considerable research and we are unaware of a complete solution. We note, however, that the estimated coefficient magnitudes are consistent with cost-share methods of estimating TFP.[14] Greenstone, Hornbeck and Moretti (2010) found that estimates drawing on this measure of TFP, using the same methodology and dataset, were similar to estimates obtained when: modeling the inputs with alternative functional forms (e.g., translog); controlling for flexible functions of investment, capital, materials, and labor; and instrumenting for current inputs with lagged changes in inputs (Syverson, 2004*a*,*b*; van Biesebroeck, 2004; Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg, Caves and Frazer, 2006; Blundell and Bond, 1998). Further, while estimated TFP can be prone to substantial measurement error, the impacts of measurement error are lessened by grouping plant-level TFP by MSA.

We estimate that MSAs' Total Factor Productivity (TFP) grew by 18%, on average, between 1977 and 1987. Note that "productivity" has a particular meaning in this context, reflecting a higher value of plant shipments for the reported levels of plant inputs. A range of factors may contribute to a place exhibiting higher TFP, including: industry mix, higher output prices or lower input prices, features of the natural environment, infrastructure or other constructed location features, or greater technological ability of manufacturing plants in the area. In the later empirical analysis, we restrict our attention to a subset of these

[12]Due to data limitations on the hours of non-production workers, total labor hours are defined as total production worker hours inflated by the ratio of all wages to production worker wages. Capital values are defined as the averaged total value of capital stocks at the beginning and end of the year, plus the total value of rentals. Capital values reflect book values, rather than the permanent inventory method, because previous annual investment data are unavailable for all plants in the Census of Manufacturers. Material inputs are defined as the total value of materials purchased minus changes in inventories.

[13]Due to Census Bureau confidentiality restrictions, MSA-year groups are omitted when they include a very small number of manufacturing firms or a very high degree of output concentration among a few firms.

[14]Consistent with constant returns to scale, the estimated coefficients on the inputs sum close to one. On average, for 1977 and 1987, the estimated coefficient is 0.58 for materials, 0.23 for labor, and 0.19 for capital. An important concern is that input choices may be adjusted endogenously in response to productivity shocks, raising the coefficient on relatively flexible inputs and attenuating the coefficient on relatively fixed inputs. The estimated impact of capital is more than one-third the combined impact of capital and labor, however, which shows little indication of attenuation in the impact of capital.

factors by instrumenting for changes in TFP with a predicted component of changes in TFP.

Of particular interest is the variation over space and time in MSA-level TFP, for which our data provide a new view drawing on aggregated plant-level TFP residuals. There is a fair degree of persistence in TFP across areas, with higher productivity places in 1980 remaining higher productivity in 1990 (Figure 1, panel A). This persistence is clearer when weighting MSAs by their manufacturing output in 1980 (Figure 1, panel B), which reflects the sample size in which TFP was estimated. Note that our reference to TFP and output data in the years 1980 and 1990 is drawn from the CMF data in years 1977 and 1987, respectively.

MSAs also experience a fair degree of relative change in TFP over time, however, despite the persistence in TFP. Figure 2 indicates the location of all sample MSAs and the size of the circle corresponds to the relative position of MSA TFP in 1980 (panel A), 1990 (panel B), and in changes from 1980 to 1990 (panel C). There are some interesting spatial patterns and there is substantial within-region variation in TFP, both in levels and in changes, which we use in the empirical analysis. How much of that variation reflects real variation in TFP, compared to how much reflects measurement error, will be an important issue to address in the empirical analysis. Most of the empirical analysis weights MSA observations by manufacturing output in 1980, where MSA-level TFP is estimated on a larger sample of manufacturing activity, and Appendix Figure 1 maps sample MSAs with the size of the circle corresponding to relative manufacturing output in 1980.

II.B Earnings, Housing Costs, and Employment

We measure labor market and housing market outcomes at the MSA level, aggregating individual-level data and household-level data from the Census of Population (IPUMS 2010). We merge TFP data from 1977 and 1987 to outcome data from 1980 and 1990, reconciling differences in the definition of MSAs across the two administrative datasets. We also include outcome data from 2000 to measure longer-run impacts of TFP, and hold constant the definition of MSAs from 1980 to 2000.

The main variables of interest are annual earnings, housing costs, and the number of workers. For some specifications, we break out the MSA-level analysis by workers' sector (manufacturing vs. non-manufacturing) or level of education (college vs. some college vs. no college). We also use data on MSAs' elasticity of housing supply from Saiz (2010).

We define a sample of adult full-time workers, following standard practice (see, e.g., Katz and Autor, 1999), by restricting the sample to: men and women between the ages of 19 and 65, who worked at least 40 weeks in the previous year, usually worked more than 35 hours per week, and worked for wages or salary in the private sector. Further, individuals' annual earnings must exceed one-half the minimum wage based on a 40-hour week and 40 weeks worked. Top-coded earnings and rents are multiplied by 1.5. The empirical estimates are not sensitive to these sample restrictions. Appendix Table 1 reports average characteristics in 1980 and average changes over time for the 193 sample MSAs.

III Econometric Models

Our goal is to estimate the impact of TFP shocks on nominal wages, employment, housing costs, and real wages. In the context of our model, we seek to estimate the parameters π_1, π_2, π_3, and π_4, as well as skill-specific versions of these parameters.

III.A Baseline Models

We regress the change in outcome Y in MSA c on the change in the MSA's average Total Factor Productivity and region dummies[15]:

$$\ln(Y_{c,t+1} - Y_{c,t}) = \pi(TFP_{c,t+1} - TFP_{c,t}) + \alpha_r + \varepsilon_c. \tag{13}$$

Our main outcomes are nominal wages, housing costs, and employment. The estimated magnitudes report the log point increase in a given outcome from a one log point increase in TFP.

Our theoretical model in Section I is based on the assumption of spatial equilibrium. In practice, it probably takes some time after a productivity shock to reach spatial equilibrium. Blanchard and Katz (1992) find regional economic adjustment within 5-10 years, though the time horizon may depend on the type of shocks under consideration. The long-run effect of a TFP shock on employment, wages, and rents will differ from its short-run effect if it takes time for workers and firms to relocate and for developers to build new housing units. In addition, in the presence of agglomeration economies, the effect of a positive local shock will tend to grow stronger over time due to self-reinforcing dynamics (Kline and Moretti, forthcoming).

We use two variants of equation 13 that allow us to distinguish short-run and long-run impacts of productivity shocks. We define short-run effects as those obtained when we use 1980-1990 changes both on the right hand side and left hand side:

$$\ln(Y_{c,1990} - Y_{c,1980}) = \pi(TFP_{c,1990} - TFP_{c,1980}) + \alpha_r + \varepsilon_c. \tag{14}$$

We do not know when the TFP shocks take place, within the 1980 to 1990 period, and the observed response from 1980 to 1990 may not reflect the full impact of a change in TFP. There may be insufficient time to fully reach the new spatial equilibrium if at least some

[15] The region fixed effects reflect the four Census regions: Northeast, Midwest, South, and West.

portion of the TFP shocks occur toward the end of that period.[16]

We define long-run effects as those obtained when we continue to measure the change in TFP between 1980 and 1990, and consider impacts on outcome changes between 1980 and 2000:

$$\ln(Y_{c,2000} - Y_{c,1980}) = \pi(TFP_{c,1990} - TFP_{c,1980}) + \alpha_r + \varepsilon_c. \tag{15}$$

This model allows for potentially delayed responses in reaching the new spatial equilibrium. Because we are interested in equilibrium outcomes, and in determining how long it takes for equilibrium to be restored, we prefer this long-run specification to a long-difference model that regresses changes from 1980 to 2000 on changes in TFP from 1980 to 2000. In that long-difference model, it continues to be unknown when the TFP shocks occur and the impacts could substantially reflect short-run effects rather than long-run equilibrium effects. In some versions of equation (15) we condition on MSAs' change in TFP from 1990 to 2000 to adjust for potential correlation in TFP shocks over time.

Elasticity of Labor Supply and Housing Supply. The reduced-form parameters (π's) are predicted to depend on the elasticity of housing supply and workers' mobility, as seen in equations 6, 7, 8, and 9. We probe the validity of our baseline estimates by examining how the estimated π's vary.

Specifically, we explore how the reduced-form parameters vary with worker mobility. We analyze data separately by worker skill groups, which we anticipate have different levels of geographic mobility, and analyze how the estimated π's vary across groups. We can also pool the specifications across skill groups, and estimate how the estimated impacts vary with worker skill s:

$$\ln(Y_{c,t+1} - Y_{c,t}) = \pi(TFP_{c,t+1} - TFP_{c,t}) + \delta(TFP_{c,t+1} - TFP_{c,t}) \times s + \alpha_{rs} + \varepsilon_c. \tag{16}$$

If the technology shock is skill-neutral, then the model predicts a city-wide productivity shock will lead to greater (relative) increases in employment of higher-skill workers and smaller (relative) increases in employment of lower-skill workers. The effect on rents depends on how integrated the housing market is across skill groups. If high and low skill workers face the same housing market within each MSA — as we assumed in our theoretical model

[16] In practice, our measured changes in TFP reflect Census of Manufacturers data from 1977 and 1987. This three-year lag period provides some additional time for the impacts to appear in the short-run estimates, though the full impact would not appear if some adjustment occurs beyond three years. Some early changes in TFP might be reflected in the 1980 data to the extent that impacts appear immediately, though in practice we expect such attenuation to be small: the adjustment period appears to be longer than a couple years (in previous literature and in our estimates), and only a small portion of the shocks would occur in the very early period.

— then the effect on rents will be identical. If high and low skill workers live in different parts of each MSA and the price shocks to one part of the market do not perfectly correlate with price shocks to the other part, then the effect on rents can be somewhat different for high and low skill workers. In practice, however, the difference cannot be too large, as the housing market would never be perfectly segmented by skill group and there is some mobility across neighborhoods.

This estimated heterogeneity is of practical importance. Variation in impacts by skill-group is informative about whether TFP shocks widen inequality or reduce inequality, particularly over the 1980's during our main period of analysis. Further, if there are indeed differences in geographic mobility, the impacts on local inequality will differ from the impacts on aggregate inequality.

We also explore how our estimates vary with local housing supply elasticity. We estimate models where the effect of a TFP shock in city c is allowed to vary depending the elasticity of housing supply in city c. To do so, we include an additional interaction term between changes in TFP and the MSA's elasticity of housing supply:

$$\ln(Y_{c,t+1} - Y_{c,t}) = \pi(TFP_{c,t+1} - TFP_{c,t}) + \delta(TFP_{c,t+1} - TFP_{c,t}) \times k_c + \alpha_r + \varepsilon_c. \tag{17}$$

Our main specifications use data from Saiz (2010) to measure the elasticity of housing supply (k_c. The theory predicts that, for a given elasticity of housing supply in other cities, a higher elasticity of local housing supply will decrease the estimated impacts on wages and housing costs.

Characterizing this heterogeneity is of practical importance in itself, beyond providing useful specification checks, because it has important policy implications for municipalities. To the extent that municipalities influence their elasticity of housing supply through land-use regulations, they could potentially influence both the local and aggregate incidence of productivity shocks. This heterogeneity also implies that the incidence of national productivity shocks will depend on where specifically those shocks occur.

III.B Biases and Instrumental Variable Models

Equations (13) and (15) present three estimation challenges.

First, estimates from these models are biased if workers are mobile, though the magnitude of this bias is likely to be trivial in our sample. Mobility introduces a bias in our reduced-form estimates because it contaminates the comparison group: as a TFP shock in one MSA draws in workers from other MSAs, a downward sloping labor demand curve would imply increased wages in other MSAs (and an upward sloping housing supply curve would imply decreased housing costs in other MSAs). Thus, the estimates of the local effect of a TFP shock on wages and rents are downward-biased and upward-biased, respectively. This bias is small, however, when each MSA is a small share of the total labor market because the general equilibrium effect is spread across many comparison MSAs. In the extreme case, where the country is of infinite size, the bias is zero. In our case, there are 325 MSAs across the United States, and so the magnitude of this bias is likely to be trivial. Aggregating these small effects across each MSA, however, may have important implications for the aggregate incidence of TFP and differential impacts by skill group — and in Section V we return to calculations of the aggregate impact.

Second, and more importantly, changes in MSA-level TFP might be correlated with unobserved factors that change the outcome variables of interest. One potentially important source of bias is represented by changes in local amenities that are both consumption and production amenities, and therefore affect both local labor demand and local labor supply. For example, improvements in the quality of local schools, community colleges, or universities could affect both the quality of the local workforce and the attractiveness of the area. These factors would both increase the local amenity level and measured TFP (since it is impossible to control perfectly for labor quality). Since better local amenities result in lower equilibrium wages and higher rents, this would lead us to underestimate the impact of TFP on wages, as higher TFP would be systematically associated with lower wages. It would also lead us to overestimate the impact on rents and employment. Similarly, an improvement in local transportation infrastructure could potentially increase both firm TFP and the desirability of the area for workers, introducing the same type of biases. Changes in local crime rates could have a similar effect, to the extent that decreased crime rates are appealing to workers and increase firm productivity. Changes in local environmental regulation may have the opposite effect: tighter environmental regulations may lower TFP, but decrease nominal wages by reducing compensating differentials associated with local environmental quality.

Third, and most practically, estimating the impacts of changes in TFP may result in substantial attenuation bias from classical measurement error in TFP. We expect this source of bias to be particularly pronounced, as TFP is notoriously difficult to measure and we use a potentially noisy estimate. Our main empirical specifications are also in first differences, and differencing the data is known to make attenuation bias more severe. Measurement error may also be non-classical, reflecting various concerns with the estimation of TFP.

To address the second and third concerns (i.e., omitted variables bias and bias from measurement error), our preferred specifications instrument for MSA-level changes in TFP using predicted MSA-level changes in TFP. By instrumenting for changes in TFP by MSA, our goal is to isolate a portion of the change in TFP across MSAs that is due to general changes in productivity and not due to particular shocks to MSAs that might otherwise be correlated with changes in MSAs' economic outcomes. In addition, predicting the change in TFP can mitigate the influence of measurement error in particular MSA's change in TFP.

We use nationwide changes in TFP by industry to predict each MSA's change in TFP, depending on which industries were initially concentrated in each MSA. Construction of the instrumental variable proceeds in three stages. First, using the plant-level data as before, we calculate residual TFP in each industry and year.[17] Second, for a particular MSA, we adjust TFP in each industry and year to reflect only plants from other MSAs. Third, for that MSA, we multiply the fraction of initial output in each industry by the national change in TFP for that industry. In this manner, the predicted change in TFP from 1977 to 1987 depends on each MSA's industries in 1977 and the change in TFP from 1977 to 1987 for those industries in other parts of the country. For example, we predict an increase in MSA-level TFP for a MSA that was initially specialized in industries that elsewhere experience an increase in TFP. We assign TFP data from 1977 and 1987 to 1980 and 1990, respectively.

The first-stage estimating equation regresses the change in TFP from 1980 to 1990 on the predicted "share-shift" change in TFP from 1980 to 1990, controlling for region fixed effects. The 2SLS estimates will reflect the estimated reduced-form impact of the predicted share-shift change in TFP on MSAs' change in outcomes, divided by this estimated first-stage impact on the actual change in MSAs' TFP. The same first-stage can be used to explore heterogeneity by the local elasticity of housing supply (equation 17) or worker type (equation 16), simply interacting the predicted change in TFP with the measure of heterogeneity.

Table 1 reports the estimated first-stage impact of the predicted share-shift change in TFP on MSAs' actual change in TFP from 1980 to 1990. In column 1, the estimated coefficient implies that a 1 log point higher predicted change in TFP is associated with a 0.80 log point higher realized change in TFP. In columns 2 and 3, we report estimates that will be used later, where we interact the change in TFP with areas' elasticity of housing supply and estimate first-stages for both endogenous variables.[18] Higher predicted TFP continues to predict higher changes in TFP, and higher TFP in more-elastic locations predicts higher changes in TFP in more-elastic locations.

We expect the IV estimates of the effect of TFP on earnings to be larger than the OLS estimates, whereas the sign of the bias is undetermined for housing costs and employment. The first and most obvious reason is the potential for classical measurement error in TFP, which would attenuate the OLS estimates, particularly when differencing the data. The second reason is the potential for changes in local amenities to affect both labor demand

[17]Industries are classified by 3-digit SIC codes, and the regressions are weighted by output.

[18]Note that housing elasticity is a fixed MSA characteristic, and so the main effect of housing elasticity is absorbed in differences.

and supply, which would downward bias OLS estimates for wages and upward bias OLS estimates for rents and employment. By contrast, predicted changes in TFP capture the portion of MSA-level TFP changes that are expected from its initial endowment of industries experiencing unrelated productivity changes, thereby avoiding bias from local sources of productivity changes (e.g., particular events in a place that might otherwise create spurious correlation between outcomes and TFP). The third reason relates to the local average treatment effect (LATE) obtained when using this instrument. Observed changes in TFP are a combination of permanent shocks and transitory shocks, such that changes in TFP are to some extent mean reverting. By contrast, our instrument may isolate variation in TFP that is more likely to be permanent because it captures industry shifts in TFP that may be long-lasting. Permanent changes in TFP are expected to have larger impacts than temporary changes in TFP, such that the IV estimates may reflect more closely the impact of a permanent change in TFP and be larger than the OLS estimates even in the absence of measurement error or unobserved changes in amenities.

Instrumenting avoids some sources of bias, though there remain some potential sources of bias. For the instrumental variables estimates, the key identification assumption is that places with different initial industry mixes would have otherwise changed similarly. The estimates would be biased upward if outcomes would have otherwise relatively *increased* in MSAs initially specializing in industries that subsequently experienced greater nationwide productivity *growth* from 1980 to 1990. It is also possible to have the reverse case, and it is unclear whether we might systematically expect the sign to go either direction. In practice, the empirical specifications control for other changes by region.

Across all specifications, we report robust standard errors adjusted for heteroskedasticity. Unless otherwise noted, the regressions are also weighted by each MSA's total manufacturing output. Our measure of TFP reflects data grouped at the MSA level, where the size of that group reflects the value of manufacturing output among sample plants. In this case of grouped data, weighting the data by group size is expected to be efficient and provides an estimate of the average impact from increasing the productivity of a fixed segment of the economy.

IV Estimates of the Local Impacts of Local TFP Shocks

IV.A Effects on Wages and Cost of Housing

Table 2 reports our baseline estimates: each panel reports estimated impacts of TFP on a particular MSA-level outcome, and each column reports estimates from a different estimating equation. Focusing on panel A, to start, column 1 reports in the cross-section that a 1 percentage point increase in TFP is associated with a 0.33% increase in annual earnings per worker. This estimated magnitude falls in column 2, estimated in changes. The estimate increases moderately from column 2 to column 3, when potentially-delayed impacts are included.

Columns 4 and 5 report estimates from our preferred specifications, which instrument for changes in TFP and report the shorter-run and longer-run impacts of TFP. The estimated magnitudes are substantially larger, consistent with the instrument mitigating attenuation bias from measurement error. In addition, omitted variables bias would downward bias the OLS estimates when higher measured TFP in particular MSAs might precisely reflect omitted variables causing workers to be paid less (panel A).

Panels B and C report estimated impacts on two proxies for housing costs: the cost of rent (panel B) and the value of owner-occupied homes (panel C). The cost of rent is perhaps the best indication of contemporaneous housing costs, though these data are only available for renters and data on home values provide a useful supplement. Higher TFP is associated with substantially higher housing costs, and the estimated effect varies across empirical specifications in a similar manner to panel A. The impact of TFP is capitalized into housing costs fairly quickly, comparing column 5 to column 4 or column 3 to column 2, though there is some delayed impact. In considering the welfare incidence of TFP, local landowners appear to gain substantially from higher local TFP.

One central economic question, relating to the incidence of TFP, is whether workers benefit from increases in TFP. Workers' earnings go up (panel A), but so do their housing costs (panel B and C). The impacts on earnings are less than the impacts on housing costs, comparing the coefficients in panels A through C, but are roughly half the impact on housing costs. Given that typical consumption shares for housing are substantially less than one-half, local workers appear to receive higher real wages from higher local TFP. Panel D reports direct estimates for real wages, under the assumption that the housing share of consumption is equal to 0.30 (Glaeser, 2008; Moretti, 2013).[19] Real wages are increasing in TFP, particularly over the long-run (column 5), as nominal wages increase by substantially more than 0.30 times the increase in housing rents.

These gains would be lessened, however, to the extent that workers also pay more for other locally-produced goods and services, whose price may adjust in response to the TFP shock. Taking the ratio of the coefficients in panels A and B, in column 5, workers would gain if they spend less than 60% of their earnings on local goods whose price increases to the same degree as housing rents.[20]

[19] In particular, we define an outcome variable equal to Log Earnings minus 0.30 times Log Cost of Rent. These estimates are similar to taking the coefficient from panel A and subtracting 0.30 times the coefficient from panel B, but here the region fixed effects are constrained to impact real wages directly.

[20] Moretti (2013) calculates a share equal to 63% for local non-tradables, inclusive of housing. Note,

The incidence of TFP on local workers will neglect an important dynamic, however, as workers are expected to move locations in response to TFP. In the cross-section, the number of workers is substantially higher in locations with greater productivity (panel D, column 1). In changes, OLS specifications find no systematic relationship between TFP and the number of workers (columns 2 and 3). The IV specifications report a substantial impact of TFP on the number of workers. Further, comparing columns 5 and 4, there is evidence of delayed labor market equilibration as additional workers flow into the MSA over the subsequent ten years.

The estimated impacts are not just statistically significant but also economically meaningful. To have a better sense of the magnitudes of our estimated coefficients, consider that TFP increased by 18% in the average city between 1980 and 1990. Our short-run estimates from column 4 indicate that, by 1990, this resulted in an increase in local earnings of 16% or $430 million for the average city (in 1980 dollars) and an increase in local rents of 33% or $119 million (assigning increases in rents to all housing units). These effects were larger in the long-run: our long-run estimates from column 5 indicate that, by 2000, there was an increase in local earnings of 26% or $698 million and an increase in local rents of 41% or $148 million.

Our model implies that the impact of TFP shocks should vary by location depending on the elasticity of housing supply. The same TFP increase should result in higher equilibrium wages and rents in cities with more inelastic housing supply than in cities with more elastic housing supply. This provides a natural test of the plausibility of our baseline model. On a substantive level, it also helps to understand whether cities can affect the incidence of TFP shocks by altering land use regulations. In other words, we can test whether more permissive land-use regulation — and therefore a higher elasticity of housing supply — changes the effect of TFP on *real wages* of residents.

Table 3 reports how the impact of TFP varies with MSAs' elasticity of housing supply. In particular, we estimate equation 17 using our preferred specifications that instrument for changes in TFP using predicted changes in TFP.[21] Column 1 reports the coefficient of interest on the interaction term: in panel A, the impact on earnings from a one log point increase in TFP is 0.099 log points smaller when MSAs' elasticity of housing supply is one standard deviation greater. Columns 2 and 3 evaluate the total effect of TFP on earnings for MSAs at the 10th and 90th percentiles of the distribution of housing elasticities: for a MSA with an inelastic housing supply (at the 10th percentile), earnings increase by 1.02%

however, that there tends to be small estimated effects of local shocks on the price of local non-tradables other than housing (Beraja, Hurst and Ospina, 2015).

[21] For these IV specifications, Table 1 previously reported the relevant first-stage estimation results.

from a 1 percentage point increase in TFP (column 2); for a MSA with an elastic housing supply (at the 90th percentile), earnings increase by 0.79%.[22]

Panels B and C report that an increase in TFP results in less of an increase in housing costs when the housing supply is more elastic. This intuitive result is in line with previous research (Saiz, 2010; Glaeser and Gyourko, 2005; Glaeser, Gyourko and Saks, 2006).

Of particular interest is how the incidence of TFP on workers varies across MSAs with different elasticities of housing supply. In more elastic locations, housing prices do not increase by as much but earnings also do not increase by as much. Panel D reports estimated increases in real earnings that are similar across MSAs with different elasticities of housing supply. Over time, there is a moderately greater impact on real wages in MSAs with a less elastic housing supply (column 4), but the incidence is positive for workers across MSAs throughout the distribution of elasticities and is broadly similar across these MSAs.

This finding, that TFP shocks impact real wages similarly in cities with high and low elasticity of housing supply, is consistent with the assumption of spatial equilibrium. For a given TFP shock and elasticity of housing supply, an equilibrium requires workers to be roughly indifferent across cities such that wages adjust proportionally to rents. On a practical level, this result suggests that the relaxation of local land-use regulations is a limited (or moderately counterproductive) mechanism through which municipalities might redistribute the economic rent generated by productivity growth from local landowners to local workers.

IV.B Differential Effects on Skilled and Unskilled Workers

Changes in local productivity may have different impacts on different types of workers. In particular, the effects can be different depending on workers' level of education. Much of the literature on the effect of technological change on wage inequality has focused on skill-biased technical change,[23] but even a skill-neutral change in local TFP might differentially impact workers with different levels of education if they have different levels of geographic mobility. The model predicts smaller local impacts on high-skill wages if high-skill workers are more geographically mobile.

We now report our baseline estimates from Table 2, but allowing for a different impact of TFP on skilled and unskilled workers. Table 4 reports that all workers' earnings increase with local TFP, but less-educated workers' earnings are particularly sensitive to local TFP (panel A, columns 1 – 3).[24] Housing costs also increase more for less-educated workers (panels B

[22] The total effect is determined by the main effect of TFP and the interaction effect between TFP and housing elasticity, evaluated at that level of housing elasticity. MSAs' housing elasticity is normalized to have a mean of zero and a standard deviation of one, and so the estimated main effects correspond to the estimates reported in Table 2.

[23] This literature is immense, though see (Acemoglu and Autor, 2011) for a recent review.

[24] Corresponding OLS estimates are reported in Appendix Table 6.

and C, columns 1 – 3), but this difference is smaller in proportional terms than the difference in impacts on earnings. Thus, panel D reports moderate increases in the real earnings of workers with no college education, and smaller and statistically insignificant impacts on real wages of more-educated workers. These estimates imply that welfare of the least-educated workers is the most dependent on local TFP, either benefiting most from an increase in TFP or losing most from a decline in TFP.

Consistent with differences in geographic mobility by workers' education level, panel E reports a moderately greater inflow of more-educated workers than less-educated workers (columns 1 – 3). These estimates are consistent with the idea that more-educated workers disproportionately leave MSAs during periods of productivity decline, and enter MSAs during times of productivity gains. These differences in worker flows exist despite greater responsiveness of earnings of low-educated workers; indeed, less-educated workers' wages may be more responsive because they are less responsive geographically.

Over time, there is greater movement of workers in response to changes in TFP and greater differences by workers' education level (Table 6, panel D, columns 5 – 8). Earnings and housing costs continue to increase for all workers, with marginally greater increases for less-educated workers (panels A – C, columns 5 – 8). Real wage increases continue to be greater for less-educated workers, but the estimated difference is smaller because more-educated workers' earnings increase by more over time.

Overall, these estimates suggest that even skill-neutral TFP increases could have greater benefits for unskilled workers in the affected city. If higher-income workers are more geographically responsive to differences in earnings, this would dissipate local increases in earnings from higher local TFP. This finding at the local level complements results in the skill-biased technical change literature, which are generally that skill-biased technological change benefits more-skilled workers at the national level. These two findings are not necessarily in contradiction, as they relate to different types of productivity shocks. Further, the geographic unit is quite different. Importantly, we have so far focused on the local effects of local shocks, and it is possible that the aggregate effects of local shocks are different, particularly as they relate to differential impacts by skill or implications for inequality. For example, while mobile high-income workers might dissipate local increases in earnings from local increases in TFP, high-income workers would benefit disproportionately in other locations. A purely local perspective may be misleading for these more aggregate implications, and we turn to aggregate impacts in Section V below.

In a complementary exercise, we look at changes in alternative measures of inequality. Table 5 reports estimated impacts of TFP on the distribution of earnings, measured as the difference in log earnings at the 90th and 10th percentiles (panel A). We then separate impacts on overall inequality into impacts on inequality within the upper portion of the distribution (panel B) and within the lower portion of the distribution (panel C). Higher TFP is associated with somewhat higher inequality in the cross-section (column 1), though this effect is largely attenuated in changes (columns 2 and 3).

In our preferred specifications, instrumenting for changes in TFP, increased TFP is associated with substantial declines in local earnings inequality (Table 5, panel A, columns 4 and 5). The estimated magnitudes imply that a 1 percentage point increase in TFP reduces the 90-10 earnings gap by 0.655 percentage points from 1980 to 1990 (i.e., earnings at the 10th percentile increase by 0.655 percentage points more than earnings at the 90th percentile). This impact on inequality occurs at the upper portions of the distribution (panels B and C).

IV.C Impacts by Sector and the Local Multiplier Effect

Increases in manufacturing TFP directly impact the manufacturing sector, but would be expected to also impact the local non-manufacturing sector. Wage and employment growth in manufacturing increase the demand for local non-traded good and services, and therefore employment and wages in non-manufacturing sectors (Moretti, 2010). Indeed, the extent to which non-manufacturing sectors are impacted is informative about how much policies directed at the manufacturing sector might influence the broader local economy.

Table 6 reports estimated effects on the manufacturing sector and non-manufacturing sectors.[25] Earnings are impacted similarly in the manufacturing sector and non-manufacturing sectors (panel A), consistent with a close integration of local labor markets.[26] If some marginal workers are able to move easily between sectors locally, then we would not expect differential impacts on earnings by sector. Panel B reports that the number of workers responds moderately more in the manufacturing sector than in the non-manufacturing sector (columns 1 – 3). These increases reflect some combination of in-migration and movement between sectors. The estimated magnitudes increase over time (columns 4 and 5), with greater increases in non-manufacturing sectors.

One empirical parameter of interest, particularly in policy settings, is the implied "multiplier effect" of the manufacturing sector on the local non-manufacturing sector. The general goal is to measure how much one additional manufacturing job creates additional local non-manufacturing jobs. Using our empirical methodology, we can calculate this multiplier using the coefficients from columns 1 and 2 of Table 6 and the average initial number of workers by sector from Appendix Table 1. From an increase in manufacturing TFP that creates one manufacturing job, panel C reports the implied increase of 1.24 non-manufacturing jobs. This

[25]Corresponding OLS estimates are reported in Appendix Table 7.

[26]Similarly, housing costs are impacted similarly for workers in the manufacturing sector and non-manufacturing sectors.

estimate is qualitatively consistent with estimates in the previous literature. For example, using a similar time horizon, (Moretti, 2010) estimates a local multiplier for manufacturing equal to 1.6.

A longer time horizon yields a larger multiplier, perhaps because it takes time for the effect of shocks in manufacturing to generate additional demand for local services. Over the long-run, based on columns 4 and 5 of Table 6, there is an implied increase of 2.96 non-manufacturing jobs. There is less empirical evidence on the long-run local multiplier, and this larger estimate suggests a larger potential for local economic incentives to generate greater local economic activity (though at the cost of displacing activity elsewhere).

IV.D Robustness

We now consider the robustness of the previous estimates to some changes in the empirical specifications.

Given the estimated changes in the number of workers, one natural question is whether changes in worker composition are driving the estimated increases in annual earnings. Appendix Table 2 reports estimated impacts on earnings after controlling for changes in worker composition, and these estimates are similar to those in Table 2. Similarly, panels B and C report similar estimates on housing costs after controlling for changes in the composition of the housing stock.

In our long-run specifications, which estimate the impact of TFP changes from 1980 to 1990 on outcome changes from 1980 to 2000, one potential concern is that long-run dynamics might be confounded with correlation over time in TFP shocks. In particular, changes in TFP between 1980 and 1990 might be correlated with changes in TFP between 1990 and 2000. Our preferred specifications instrument for changes in TFP, and so predicted changes in TFP between 1980 and 1990 may be less correlated with actual changes in TFP between 1990 and 2000. Appendix Table 3, column 1, reports similar estimates (to those in column 5 of Table 2) when controlling for the measured change in TFP between 1990 and 2000. Column 2, also reports similar estimates when controlling for the measured change in TFP between 1990 and 2000 and instrumenting for the measured change with the predicted change in TFP between 1990 and 2000. Column 3 reports somewhat smaller estimates from a long difference specification, regressing outcome changes from 1980 to 2000 on TFP changes from 1980 to 2000, and instrumenting with the predicted change in TFP from 1980 to 2000. This long difference specification need not reflect long-run effects, however, as changes in TFP could occur any time between 1980 and 2000.[27]

[27] There is a less robust first-stage relationship between predicted TFP changes from 1990 to 2000 and actual TFP changes from 1990 to 2000, which contributes to our focus on changes in TFP from 1980 to 1990.

Appendix Table 4 reproduces estimates from Table 2, but weights the regressions by MSA population in 1980 (panel A) or does not weight the regressions (panel B) rather than weighting by MSA's manufacturing output in 1980. The interpretation of the results is broadly similar, with somewhat smaller magnitudes in the unweighted 2SLS specifications. The OLS results are substantially attenuated in the unweighted specifications, particularly when estimated in changes, which is consistent with measurement error in TFP that is mitigated by focusing on MSAs with greater manufacturing output and more-precisely measured TFP.

In considering the differential effects of TFP by MSAs' housing elasticity, Table 3 had only included the 2SLS results. Appendix Table 5, panel A, reports interaction effects for the cross-section (column 1), the OLS estimates in changes (column 2), and the 2SLS estimates in changes as a basis for comparison (columns 4 and 5). The estimated interaction effects are quite similar in the OLS specifications, both in the cross-section and in changes. Panel B reports similar estimated interaction effects when controlling for interactions between region fixed effects and changes in TFP, which focus the regression on within-region variation in housing elasticity. The distribution of housing elasticities is fairly right-skewed at the extremes (toward high elasticity locations), but panel C reports similar estimates when omitting MSAs with the 10 largest and 10 smallest elasticities of housing supply. Panel D reports estimates that are similar, with somewhat larger magnitudes, when using an alternative measure of MSAs' elasticity of housing supply (Gyourko, Saiz and Summers, 2008).[28]

Appendix Tables 6 and 7 report OLS estimates of TFP impacts by worker skill group and sector, as a basis for comparison, that correspond to the 2SLS estimates reported in Tables 4 and 6. For reasons discussed above, however, we focus on the 2SLS specifications.

V Estimates of the Aggregate Impacts of Local TFP Shocks

The estimates in Section IV indicate that local TFP shocks increase local wages in the metropolitan area where the shocks occur, relative to other metropolitan areas. TFP shocks also increase local housing costs, but to a lesser extent such that local real wages also increase.

These direct impacts on local outcomes are only part of the overall impact from local TFP shocks, however. Our model indicates that a positive shock to a city's TFP has impacts on earnings and housing costs in other MSAs through labor markets that are partially integrated across geographic space.

[28]Our main measure of housing elasticity, from Saiz (2010), combines information on geographic housing constraints on housing supply with this measure of housing elasticity that reflects regulatory constraints on housing supply. There is a large literature on housing supply (see, e.g., Glaeser, Gyourko and Saks, 2005, 2006; Glaeser and Ward, 2009), and Gyourko (2009) provides a recent survey.

In this section, we quantify the aggregate impact of local TFP shocks on earnings and housing costs, distinguishing between the *direct* impact on the cities experiencing the shocks and the *indirect* impact on other cities. We previously estimated substantial geographic mobility in response to changes in TFP, particularly in the long-run, and so we expect there to be potentially substantial indirect impacts of TFP shocks.

V.A Aggregate Direct Effect

We begin by estimating the sum of all the localized long-run effects on the cities *directly* affected by TFP shocks between 1980 and 1990.

We start with the empirical distribution of changes in TFP from 1980 to 1990. For each MSA, we use our estimates in Table 3 to quantify the aggregate impact of these observed TFP shocks on overall earnings and housing costs in the long-run on the cities where the shocks took place. Geography matters in our setting, in the sense that the same TFP shock has potentially different aggregate impacts depending on where it occurs because of MSA-level differences in the elasticity of housing supply. As shown in Table 3, a 1 percentage point increase in TFP in a city with a more inelastic housing supply results in a significantly larger increase in earnings than a 1 percentage point increase in TFP in a city with a more elastic housing supply.

In practice, for each MSA, we multiply the observed TFP shock by our preferred estimate of the long run impact of TFP on earnings and scale it by the initial level of earnings in 1980. In particular, we use the long-run instrumental variable estimates for the main effect of TFP and the interaction with housing elasticity (i.e., estimates from column 4 of Table 3 and the corresponding main effect, which are both reflected in columns 5 and 6 of Table 3). We scale these magnitudes to the national level, multiplying by the ratio of national annual wage earnings to the wage earnings of sample workers in our sample MSAs, and convert 1980 dollars to 2014 dollars using the CPI.

For housing costs, we multiply annual housing costs in each MSA in 1980 by the corresponding estimate of the impact of TFP on housing rents. We scale these magnitudes to the national level, multiplying by the ratio of national housing value to the value of housing for sample workers in our sample MSAs, and convert 1980 dollars to 2014 dollars using the CPI.[29]

Table 7, panel A and column 1, reports a direct impact on total annual earnings of $989 billion from increases in TFP between 1980 and 1990. Column 2 reports a direct impact

[29]To obtain total annual housing costs for our sample workers, we take the number of unique sample households in each MSA and multiply by the average annual rent of renter-occupied housing. To scale this number to the national level, where the rental market may be different, we use the ratio of total owner-occupied house value nationally and total owner-occupied house value in the sample.

on total housing costs of $245 billion. These estimates are simply a scaling of our previous regression estimates, but provide a benchmark for the following calculations.

V.B Aggregate Indirect Effect

We now turn to the effect of local TFP shocks on the cities not directly affected. For each city c directly affected by a TFP shock between 1980 and 1990, we estimate the impact on wages and housing costs on all other cities, and then sum these indirect effects across all cities c directly affected.

As before, geography can matter because of heterogeneity in the elasticity of housing supply. Thus, the same TFP shock has potentially different indirect effects depending on elasticity of housing supply in the city directly hit by the shock (which affects the number of movers) and the elasticity of housing supply in all other cities (which affects the magnitude of indirect wage and housing cost effects).

We proceed in three steps. First, for each city c, we calculate the number of workers drawn to that city.[30] Second, we calculate the associated number of workers leaving each other MSA based on different assumptions regarding migratory flows (discussed below).[31] Third, depending on the change in employment in each city indirectly affected, we calculate the resulting change in wages (based on the assumed elasticity of labor demand) and the resulting change in housing costs (based on the measured elasticity of housing supply). Fourth, we sum the effects from one MSA on all other MSAs. Fifth, and finally, we sum these aggregate effects from all MSAs.

The above calculations require two additional assumptions. The first additional assumption concerns equilibrium worker flows; specifically, how many of the workers attracted to city c are coming from each other city. As a benchmark scenario, we assume that workers moving to a particular city are drawn from all other cities in proportion to their number of workers. In this scenario, worker migration does not depend on geographic distance, but it is a useful benchmark that holds fixed the relative sizes of other cities. At the same time, this scenario is not very realistic because mobility patterns vary across the country. As an alternative scenario, we assume that workers moving to a particular city are drawn from all other cities in proportion to actual city-to-city flows observed between 1975 and 1980. These years have the advantage of predating our sample period, whereas mobility flows in

[30] This number of workers depends on that city's TFP shock between 1980 and 1990, that city's elasticity of housing supply, and the 2SLS estimated long-run impact on employment (as estimated in column 5 of Table 2, but extended to allow for an interaction with housing elasticity). In practice, we do not find substantially fewer movers to cities with a more inelastic supply of housing.

[31] For these aggregate calculations only, when MSAs in our main sample are missing data on the elasticity of housing supply, we assign an elasticity of housing supply to those MSAs that is equal to the average of other MSAs in the same state (or the same Census region if there are no other MSAs with housing elasticity data in the same state).

later years could be endogenous to changes in TFP. City-to-city flows reflect variation in geographic and economic distance between cities, though direct migration flows may not reflect the eventual redistribution of population across all cities (whereas our benchmark assumption holds fixed the relative sizes of other cities, which are generally persistent over time). In both cases we assume a closed economy without international migration, in which a fixed number of workers move across cities.

The second additional assumption concerns the production technology across cities. To calculate the effect on wages from given labor supply shifts in each city indirectly affected by the shock, we need to assume a value for the elasticity of labor demand in each city. In our baseline estimates, we assume a labor share of 0.65, which corresponds to assuming that parameter h is equal to 0.35 in equation 1, and implies that a 1% decline in the labor force will increase wages by 0.35%. The assumed labor share of 65 percent is consistent with calculations by Karabarbounis and Neiman (2014).

Summing the indirect impacts generated by worker movement, Table 7 (panel A) reports that the direct effects in columns 1 and 2 were associated with an indirect \$845 billion increase in earnings (column 3) and an indirect \$328 billion decrease in housing costs (column 4). That is, the indirect effects contribute an additional 85% to the direct effect on earnings and roughly counterbalance the direct effects on housing costs.[32] Summing the direct and indirect effects, TFP growth between 1980 and 1990 is calculated to have raised earnings by \$1.83 trillon between 1980 and 2000 (in constant 2014 dollars).

Panel B reports that the indirect effects are moderately smaller when assigning MSA-to-MSA worker flows based on 1975-1980 flows: an indirect \$579 billion increase in earnings and an indirect \$207 billion decrease in housing costs. The indirect effects contribute an additional 59% to the direct effect on earnings and continue to roughly counterbalance the direct effects on housing costs.

Panel C reports similar results for renters only. Local and aggregate incidence is clearest for renters, who are not also affected through their ownership of housing.

Panels D, E, and F split the analysis by MSAs' elasticity of housing supply. Panel D reports direct and indirect impacts from increases in TFP among the MSAs with the least elastic housing supply (i.e., those in the bottom 10% of the distribution); panel E corresponds to MSAs with moderate elasticities of housing supply; and panel F corresponds to MSAs with the most elastic housing supply (i.e., those in the top 10% of the distribution).[33] The

[32]The sum of direct and indirect impacts on housing costs is close to zero because we have assumed a closed economy with a fixed number of workers. Note that the total impact on housing costs is not identical to zero because the elasticity of housing supply differs across cities.

[33]Note that we report indirect impacts on all other MSAs from increases in TFP among the indicated subset of MSAs.

absolute magnitudes are very different across panels, which reflects differences in the size and number of MSAs, but our focus is on the comparison across columns within each panel. Less-elastic MSAs generate less of an indirect increase in earnings, relative to the direct effect on earnings, because there is less movement of workers in response to changes in TFP. Similarly, less-elastic MSAs generate less of an indirect decrease in housing costs, compared to the direct increase in housing costs, because workers are being drawn from MSAs with a more elastic supply of housing.

In the aggregate, the incidence of TFP shocks falls mainly on workers. Impacts on landowners in one MSA are counterbalanced by impacts on landowners in other MSAs. Further, while local gains to workers are dissipated by the inflow of additional workers, this generates gains to workers in other locations with increasing labor scarcity.

V.C Direct and Indirect Effects by Educational Level

Aggregate economic effects may differ importantly by workers' education levels, as spillover effects arise due to worker mobility and the previous reduced-form estimates showed differences across education levels in the responsiveness of workers to local changes in TFP. Table 8 reports estimated direct and indirect effects on wages and housing costs for all workers again (in panel A), but then also for workers with no college education (panel B) and workers with a college degree (panel C).

In these calculations, the geography of TFP shocks matters for two reasons. First, as before, the same shocks have potentially different effects due to heterogeneity in the elasticity of housing supply. Second, and in addition to before, the same shocks have potentially different effects due to variation across cities in their share of high-skilled workers.

We find that the spillover effects on wages are a greater share of the total effect on wages for higher skilled workers (column 3 as a fraction of columns 1 and 3). Similarly, the spillover effects on housing costs are larger and more negative than the direct effects on housing costs for higher skilled workers (column 4 compared to column 2). These estimates imply that adjusting for spillover effects has the consequence of raising the incidence on more-educated workers relative to less-educated workers.

Because high education workers move in greater numbers in response to changes in local TFP, these TFP shocks generate greater positive spillover effects on high education workers elsewhere. These spillover effects partly counterbalance the lower local incidence of TFP shocks on high education workers (Table 4). Overall, however, when inflating the wage estimates from Table 4 to adjust for the magnitude of spillover effect (from the ratio of column 3 to column 1 in Table 8), the percent increase in wages continues to be greater for workers with less education.

VI Conclusion

Our estimates suggest three main conclusions.

First, while local productivity growth benefits local landowners, there are also direct benefits for workers in both nominal and real terms. Between 1980 and 1990, TFP growth increased nominal earnings of sample workers by $430 million and local rents by $119 million. In this respect, local real earnings increased by $311 million (in 1980 dollars). These increases were larger in the long-run: by 2000, local real earnings had increased by $550 million.

Second, the local effects of even factor-neutral TFP shocks may differ importantly for skilled and unskilled workers. Consistent with lower geographic mobility among less-educated workers, we estimate that increases in TFP are associated with greater percent increases in the number of college graduates than in the number of high school graduates. These estimates imply that the local supply of high school graduates is less elastic to changes in TFP than that of college graduates. Following a shift in labor demand, due to increased TFP, skill groups with a less elastic labor supply should receive larger wage effects. Indeed, we estimate that local increases in TFP resulted in larger wage increases for high school graduates than for college graduates. As a consequence, local increases in TFP compressed inequality at the local level. In practice, the observed TFP shocks between 1980 and 1990 may have been skill-biased or not, but these estimates highlight how differences in geographic mobility can generate skill-biased impacts from even factor-neutral TFP shocks.

Third, a large fraction of the overall impact of local TFP shocks are not localized. In particular, almost one-half of the overall increase in workers' nominal earnings occurs in cities not directly impacted by TFP changes. There is little aggregate impact on housing costs, as the increases in cities directly impacted by TFP gains are mostly offset by losses in other cities. Importantly, the indirect economic impacts are different for high-skill and low-skill workers. As a consequence, the aggregate incidence of TFP shocks differs from the local incidence, and in a manner more skewed toward high-skill workers.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." , ed. O. Ashenfelter and D. Card Vol. 4 of *Handbook of Labor Economics*, Chapter 12, 1043–1171. Elsevier.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2006. "Structural Identification of Production Functions." *UCLA Mimeo.*

Bentolila, Samuel, and Gilles Saint-Paul. 2003. "Explaining Movements in the Labor Share." *Contributions to Macroeconomics*, 3(1): 1–33.

Beraja, Martin, Erik Hurst, and Juan Ospina. 2015. "The Aggregate Implications of Regional Business Cycles." Mimeo.

Blanchard, Oliver J. 1997. "The Medium Run." *Brookings Papers on Economic Activity*, 28(2): 89–158.

Blanchard, Olivier J., and Francesco Giavazzi. 2003. "Macroeconomic Effects of Regulation and Deregulation in Goods and Labor Markets." *Quarterly Journal of Economics*, 118(3): pp. 879–907.

Blanchard, Olivier J., and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity*, 1992(1): pp. 1–75.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics*, 87(1): 115–143.

Caseli, Francesco. 1999. "Accounting for Cross Country Income Differences." *Handbook of Economic Growth*, , ed. Philippe Aghion and Stephen Durlauf, 1463–1555. North-Holland: Elsevier.

Glaeser, Edward L. 2008. *Cities, Agglomeration and Spatial Equilibrium.* Oxford University Press.

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): 265–278.

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy*, 113(2): pp. 345–000.

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., Joseph Gyourko, and Raven Saks. 2005. "Why is Manhattan so Expensive? Regulation and the Rise in Housing Prices." *Journal of Law and Economics*, 48.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy*, 118(3): 536 – 598.

Gyourko, Joseph. 2009. "Housing supply." *Annual Review of Economics*, 1(1): 295–318.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies*, 45(3): 693–729.

Hall, Robert E., and Charles I. Jones. 1999. "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics*, 114(1): pp. 83–116.

Hsieh, Chang-Tai, and Enrico Moretti. 2015. "Why Do Cities Matter? Local Growth and Aggregate Growth." Mimeo.

Jones, Charles I. 2005. "The Shape of Production Functions and the Direction of Technical Change." *Quarterly Journal of Economics*, 120(2): 517–549.

Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics*, 129(1): 61–103.

Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." *Handbook of Labor Economics, vol. 3A*, , ed. O. Ashenfelter and D. Card, 1463–1555.

Kline, Patrick, and Enrico Moretti. forthcoming. "Local Economic Development, Agglomeration Economies and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority." *Quarterly Journal of Economics*.

Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies*, 70(2): 317–342.

Malamud, Ofer, and Abbigail Wozniak. 2012. "The Impact of College Education on Geographic Mobility." *Journal of Human Resources*, 47(4): 913–950.

Moretti, Enrico. 2010. "Local Multipliers." *American Economic Review*, 100(2): 373–77.

Moretti, Enrico. 2011. "Local labor markets." *Handbook of Labor Economics*, 4: 1237–1313.

Moretti, Enrico. 2013. "Real Wage Inequality." *American Economic Journal: Applied Economics*, 5(1): 65–103.

Olley, Steven G., and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263–1298.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century.* Harvard University Press.

Rodriguez, Francisco, and Arjun Jayadev. 2010. "The Declining Labor Share of Income." Human Development Report Office (HDRO), United Nations Development Programme (UNDP) Human Development Research Papers (2009 to present) HDRP-2010-36.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics*, 125(3): 1253–1296.

Syverson, Chad. 2004*a*. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6): 1181–1222.

Syverson, Chad. 2004*b*. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics*, 86(2): 534–550.

van Biesebroeck, Johannes. 2004. "Robustness of Productivity Estimates." *NBER Working Paper 10303*.

Wozniak, Abigail. 2010. "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources*, 45(4): 944–970.

Figure 1. Total Factor Productivity (TFP) by MSA, 1980 and 1990

Panel A. Unweighted

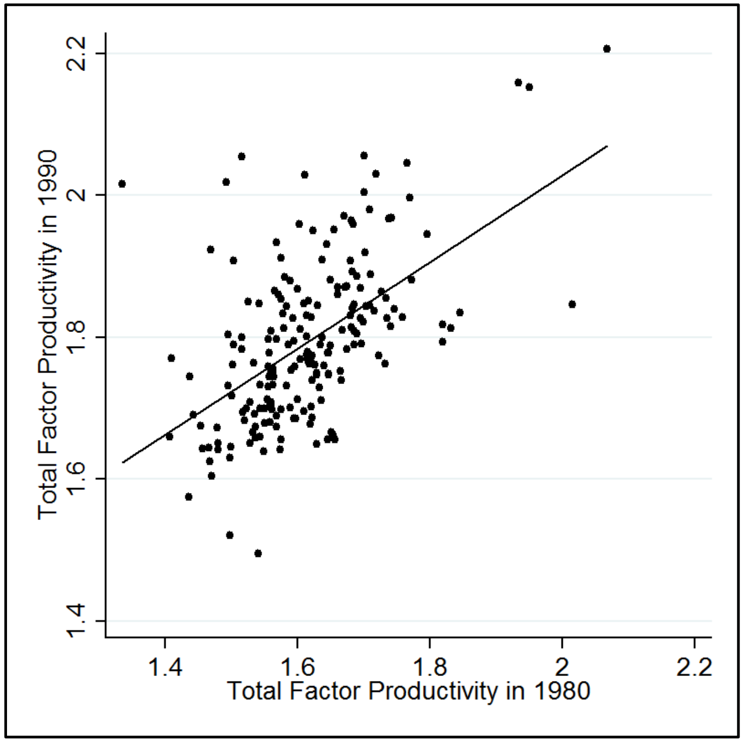

Panel B. Weighted by Manufacturing Output in 1980

Notes: Panel A plots TFP in 1990 against TFP in 1980 for each MSA in the sample along with the best linear fit for the relationship. Panel B plots TFP in 1990 against TFP in 1980, where the dot size reflects MSA manufacturing output in 1980. The linear fit in Panel B is also weighted by MSA manufacturing output.

Figure 2. Spatial Distribution of Total Factor Productivity, 1980 and 1990

Panel A. TFP in 1980

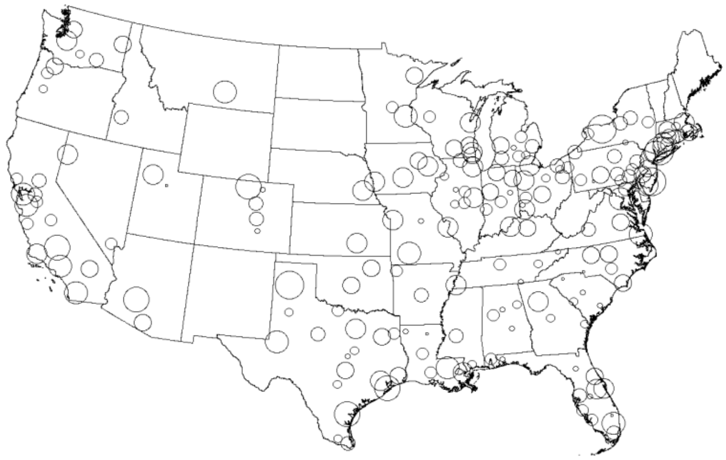

Panel B. TFP in 1990

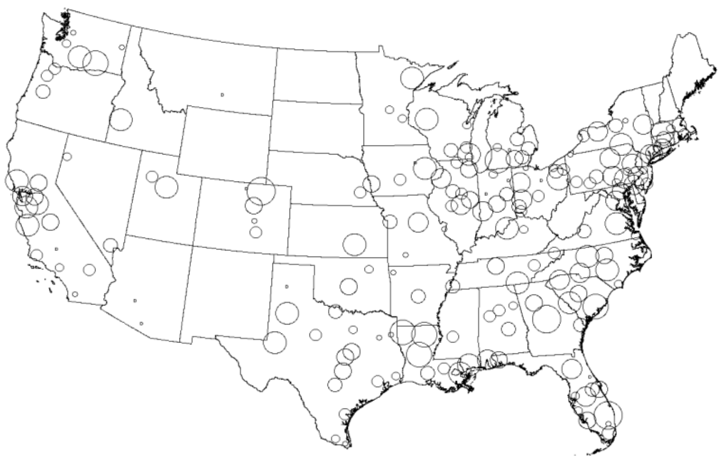

Panel C. Change in TFP from 1980 to 1990

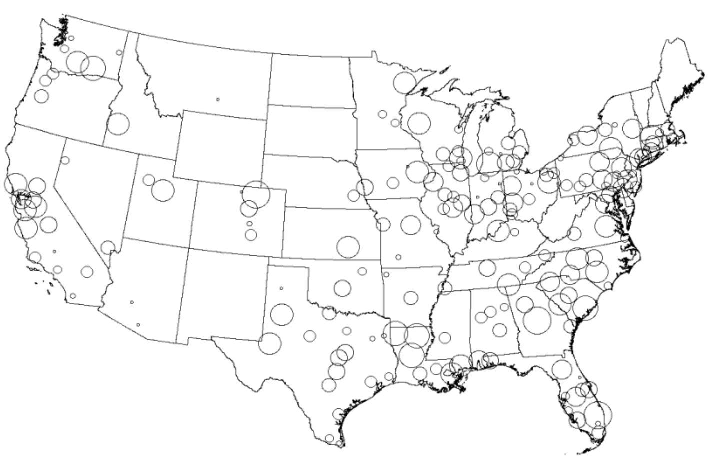

Notes: Panels A and B map total factor productivity (TFP) in 1980 and 1990. Each circle corresponds to a MSA, and the size of the circle reflects that MSA's relative TFP. Panel C depicts the change in total factor productivity (TFP) from 1980 to 1990 for each MSA in the sample. The relative size of each circle corresponds to the relative change in TFP (i.e., larger circles reflect a relative increase in TFP).

Table 1. First-Stage: Change in TFP and Predicted Change in TFP

	Change in TFP, 1980 to 1990	Change in TFP, 1980 to 1990	Change in TFP, Interacted with Elasticity
	(1)	(2)	
Predicted Change in TFP, 1980 to 1990	0.80***	0.82***	-0.02
	(0.17)	(0.17)	(0.17)
Predicted Change in TFP, 1980 to 1990, interacted with Elasticity		-0.03	1.31***
		(0.07)	(0.12)
Sample MSAs	193	168	

Notes: Column 1 reports the first-stage coefficient from a regression of county TFP on the predicted change in TFP, calculated from each MSA's distribution of industries in 1980 and changes in all other MSAs in TFP by industry from 1980 to 1990. The regression is weighted by county manufacturing output and we include census region fixed effects. Column 2 reports the same specification as in column 1, but TFP is also interacted with each MSA's elasticity of housing supply. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 2. Estimated Impacts of TFP on Earnings, Housing Costs, and Employment

	Cross-section in 1980 and 1990 OLS (1)	Change from 1980 to 1990 OLS (2)	Change from 1980 to 2000 OLS (3)	Change from 1980 to 1990 2SLS (4)	Change from 1980 to 2000 2SLS (5)
Panel A. Log Earnings	0.33***	0.14*	0.29**	0.91***	1.45***
	(0.07)	(0.07)	(0.12)	(0.32)	(0.47)
Panel B. Log Cost of Rent	0.83***	0.32**	0.46***	1.86***	2.30***
	(0.15)	(0.15)	(0.15)	(0.66)	(0.66)
Panel C. Log Home Value	1.10***	0.48**	0.69***	1.74**	2.46***
	(0.24)	(0.23)	(0.20)	(0.72)	(0.78)
Panel D. Log Real Earnings	0.08*	0.04	0.15*	0.35*	0.76**
	(0.05)	(0.05)	(0.09)	(0.19)	(0.32)
Panel E. Log Workers	3.59***	0.05	0.15	2.37***	4.16***
	(1.06)	(0.17)	(0.23)	(0.80)	(1.26)

Notes: Column 1 reports estimates from equation 17 in the text: the indicated MSA characteristic from each panel is regressed on MSA total factor productivity (TFP) in 1990 and 1980, controlling for region-by-year fixed effects and weighting each MSA by its total manufacturing output. Column 2 reports estimates from equation 19 in the text: the indicated MSA characteristic from each panel (in changes from 1980 to 1990) is regressed on the change in MSA TFP from 1980 to 1990, controlling for region fixed effects and weighting each MSA by its total manufacturing output in 1980. Column 3 reports estimates from equation 20 in the text: the indicated MSA characteristic from each panel (in changes from 1980 to 2000) is regressed on the change in MSA TFP from 1980 to 1990, controlling for region fixed effects and weighting each MSA by its total manufacturing output in 1980. Columns 4 and 5 report 2SLS estimates corresponding to columns 2 and 3, but instrumenting for the change in TFP using the predicted change in TFP (as described in the text and notes for Table 1). The corresponding first-stage estimate is reported in Table 1. In Panel D, Log Real Earnings are defined as: Log(Earnings) - 0.3*Log(Cost of Rent). Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 3. Variation in the Estimated Impacts of TFP, by MSAs' Elasticity of Housing Supply

	Short-run Effect: Change from 1980 to 1990 (2SLS)			Long-run Effect: Change from 1980 to 2000 (2SLS)		
	Interaction Effect	Predicted Impact at 10th Percentile	Predicted Impact at 90th Percentile	Interaction Effect	Predicted Impact at 10th Percentile	Predicted Impact at 90th Percentile
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Log Earnings	-0.099**	1.02***	0.79**	-0.203***	1.69***	1.21***
	(0.048)	(0.31)	(0.31)	(0.070)	(0.44)	(0.45)
Panel B. Log Cost of Rent	-0.372***	2.27***	1.40**	-0.298***	2.64***	1.95***
	(0.089)	(0.63)	(0.64)	(0.094)	(0.63)	(0.66)
Panel C. Log Home Value	-0.478***	2.26***	1.14*	-0.400***	2.87***	1.93***
	(0.109)	(0.68)	(0.68)	(0.141)	(0.69)	(0.71)
Panel D. Log Real Earnings	0.013	0.34*	0.37**	-0.114**	0.90***	0.63**
	(0.027)	(0.19)	(0.19)	(0.049)	(0.31)	(0.31)

Notes: Column 1 reports the coefficient on the interaction term from estimating equation 21 in the text: the indicated MSA characteristic from each panel (in changes from 1980 to 1990) is regressed on the change in MSA TFP from 1980 to 1990, and the change in TFP interacted with the MSA's elasticity of housing supply (normalized to have a standard deviation of one). The first-stage estimates are reported in Table 1, where the change in TFP (main effect and interaction effect) is instrumented using the predicted change in TFP (alone and interacted with the MSA's elasticity of housing supply). Otherwise, the specification is the same as in column 4 of Table 2: including region fixed effects and weighting each MSA by its total manufacturing output in 1980. Based on the estimated main effect of TFP and the interaction effect between TFP and the elasticity of housing supply, Columns 2 and 3 report the implied effect of a one log point increase in TFP for MSAs at the 10th centile of the elasticity distribution and at the 90th centile of the elasticity distribution, respectively. Columns 4 -- 6 report the same estimates as in columns 1 -- 3, but when the outcome variable is defined to be the change in the MSA characteristic from 1980 to 2000 (indicated for each panel). Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 4. Estimated Impacts of TFP, by Education Level

	Short-run Effect: Change from 1980 to 1990 (2SLS)			Difference	Long-run Effect: Change from 1980 to 2000 (2SLS)			Difference
	College	Some College	No College	(1) - (3)	College	Some College	No College	(5) - (7)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Log Earnings	0.60**	0.67***	1.12***	-0.52***	0.87**	1.13***	1.23***	-0.36
	(0.24)	(0.26)	(0.30)	(0.20)	(0.34)	(0.33)	(0.35)	(0.32)
Panel B. Log Cost of Rent	1.13*	1.85***	2.04***	-0.91***	1.56**	2.22***	2.43***	-0.87**
	(0.61)	(0.59)	(0.71)	(0.29)	(0.61)	(0.56)	(0.70)	(0.41)
Panel C. Log Home Value	1.58***	1.69**	1.99***	-0.40	1.83***	2.19***	2.62***	-0.79**
	(0.58)	(0.74)	(0.77)	(0.30)	(0.59)	(0.74)	(0.77)	(0.32)
Panel D. Log Real Earnings	0.26	0.11	0.51***	-0.25	0.40*	0.46**	0.50***	-0.10
	(0.18)	(0.13)	(0.14)	(0.21)	(0.23)	(0.20)	(0.19)	(0.27)
Panel E. Log Workers	2.79**	2.59***	2.31***	0.48	5.82***	4.74***	3.32***	2.50**
	(1.13)	(0.73)	(0.78)	(0.66)	(1.87)	(1.24)	(1.16)	(1.21)

Notes: Columns 1 - 3 report estimates that correspond to those in column 4 of Table 2, but separately by skill group: those with a college degree (column 1), those with some college education (column 2), and those with no college education (column 3). Column 4 reports the difference between column 1 and column 3. Columns 5 - 8 report analogous estimates for the long-run effect by skill-group, corresponding to the estimates in column 5 of Table 3. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 5. Estimated Impacts of TFP on Earnings Inequality

	Cross-section in 1980 and 1990 OLS (1)	Change from 1980 to 1990 OLS (2)	Change from 1980 to 2000 OLS (3)	Change from 1980 to 1990 2SLS (4)	Change from 1980 to 2000 2SLS (5)
Panel A. Log 90/10 Centile	0.139**	-0.040	0.068	-0.655***	-1.006**
Earnings	(0.054)	(0.068)	(0.131)	(0.236)	(0.427)
Panel B. Log 90/50 Centile	0.139***	-0.074*	-0.098	-0.571***	-0.919***
Earnings	(0.048)	(0.045)	(0.061)	(0.221)	(0.316)
Panel C. Log 50/10 Centile	-0.000	0.034	0.166	-0.085	-0.088
Earnings	(0.044)	(0.060)	(0.103)	(0.240)	(0.295)

Notes: Each column reports estimates analogous to those reported in Table 2, but for MSA-level outcomes that correspond to earnings inequality: the difference between log earnings at the 90th centile and the 10th centile of the MSA's earnings distribution (Panel A), the difference between log earnings at the 90th centile and the 50th centile (Panel B), and the difference between log earnings at the 50th centile and the 10th centile (Panel C). Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 6. Estimated Impacts of TFP, by Sector

	Short-run Effect: Change from 1980 to 1990 (2SLS)		Difference	Long-run Effect: Change from 1980 to 2000 (2SLS)		Difference
	Manufacturing	Non-Manufacturing	(1) - (2)	Manufacturing	Non-Manufacturing	(4) - (5)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Log Earnings	0.75**	0.83***	-0.09	0.87**	1.45***	-0.58**
	(0.30)	(0.29)	(0.18)	(0.37)	(0.46)	(0.27)
Panel B. Log Employment	2.60***	2.16***	0.44	3.75***	4.13***	-0.38
	(0.95)	(0.70)	(0.41)	(1.26)	(1.17)	(0.48)
Panel C. Implied Multiplier	1.24**			2.96***		
	(0.53)			(1.45)		

Notes: In Panels A and B, columns 1 and 2 report estimates that correspond to those in column 4 of Table 2, but separately for the manufacturing sector (column 1) and non-manufacturing sectors (column 2). Column 3 reports the difference between column 1 and column 2. Columns 4 - 6 report analogous estimates for the long-run effect by sector, corresponding to the estimates in column 5 of Table 2. Panel C reports the implied multiplier effect: the number of increased jobs in non-manufacturing sectors associated with a increase of one job in the manufacturing sector. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 7. Aggregate Direct and Indirect Effects From Increased TFP (2014 USD, Billions)

	Direct Effect on:		Indirect Effect on:	
	Annual Earnings	Housing Costs	Annual Earnings	Housing Costs
	(1)	(2)	(3)	(4)
Panel A. All Workers	989	245	845	-328
	(274)	(61)		
Panel B. All Workers, 1975-1980 Flows	989	245	579	-207
	(274)	(61)		
Panel C. Renters Only	258	54	212	-116
	(72)	(21)		
Panel D. Low Elasticity MSAs	429	114	308	-116
(Bottom 10%)	(111)	(26)		
Panel E. Middle Elasticity MSAs	538	126	505	-200
(Middle 80%)	(156)	(33)		
Panel F. High Elasticity MSAs	22	5	31	-12
(Highest 10%)	(9)	(2)		

Notes: Panel A reports the estimated direct effects (columns 1 and 2) and indirect effects (columns 3 and 4) of MSA-level changes in TFP from 1980 to 1990 on changes from 1980 to 2000 in national earnings (columns 1 and 3) and housing costs (columns 2 and 4). As described in the text, these numbers are scaled to the national level and adjusted to billions of US dollars in 2014. Panel B reports estimates when assuming that employment displacement follows MSA-to-MSA migration flows between 1975 and 1980 (rather than being proportional to MSA populations). Panel C reports estimates for renters only. Panel D reports estimates for TFP changes in MSAs with the lowest elasticity of housing supply, Panel E reports estimates for TFP changes in MSAs with typical elasticities of housing supply, and Panel F reports estimates for TFP changes in MSAs with the highest elasticity of housing supply. Robust standard errors are reported in parentheses.

Table 8. Aggregate Direct and Indirect Effects From Increased TFP (2014 USD, Billions)

	Direct Effect on:		Indirect Effect on:	
	Annual Earnings	Housing Costs	Annual Earnings	Housing Costs
	(1)	(2)	(3)	(4)
Panel A. All Workers	989	245	845	-328
	(274)	(61)		
Panel B. Workers with No College	388	125	346	-137
	(112)	(33)		
Panel C. Workers with College Degree	163	49	329	-113
	(54)	(16)		

Notes: Panel A reports the same estimates as in Panel A of Table 7. Panels B and C report corresponding estimates for workers with no college education and for workers with a college degree, respectively. Robust standard errors are reported in parentheses.

Appendix Figure 1. Spatial Distribution of Total Manufacturing Output by MSA

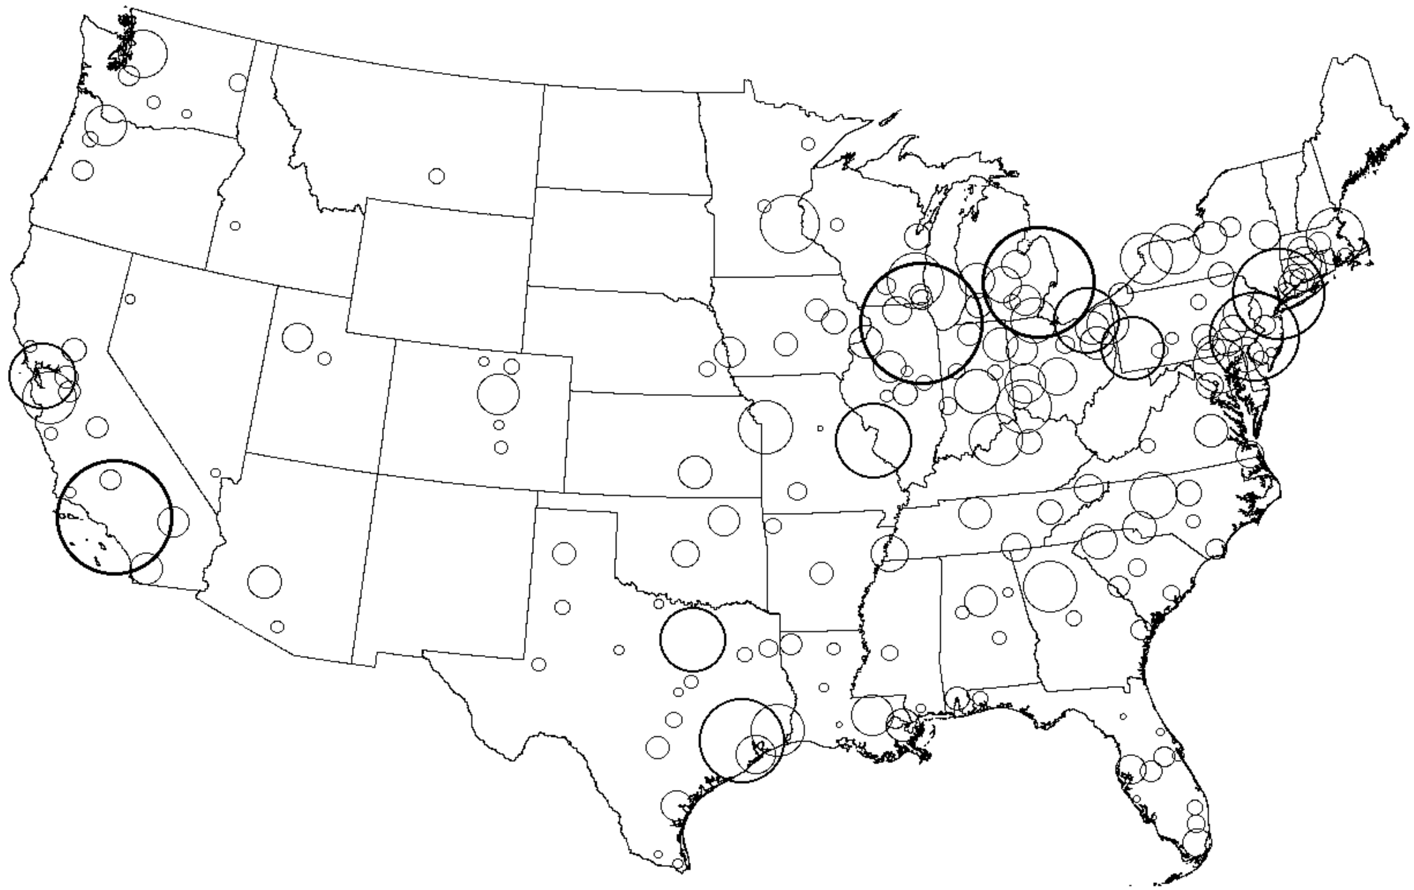

Notes: This map depicts manufacturing output for each MSA in the sample for 1980. Relatively large manufacturing output is represented by large circles.

Appendix Table 1. Baseline MSA Characteristics and Average Changes Over Time

	MSA Mean in:	Log Change in MSA Mean from:	
	1980	1980 to 1990	1980 to 2000
MSA Characteristic:	(1)	(2)	(3)
Annual Earnings	15412	0.520	0.892
	[1799]	[0.081]	[0.118]
Annual Cost of Rent	2639	0.577	0.890
	[537]	[0.159]	[0.133]
Home Value	55814	0.462	0.874
	[17428]	[0.254]	[0.184]
Number of Workers	174217	0.119	0.392
	[355645]	[0.234]	[0.319]
Workers, Completed College	31712	0.287	0.710
	[74470]	[0.285]	[0.362]
Workers, Some College	36276	0.590	0.915
	[74466]	[0.221]	[0.296]
Workers, Completed High School	106230	-0.154	0.017
	[209268]	[0.238]	[0.334]
Workers, Manufacturing Sector	57870	-0.030	0.091
	[120461]	[0.294]	[0.410]
Workers, Non-Manufacturing	116347	0.197	0.524
	[239858]	[0.220]	[0.292]
Number of Housing Units	137195	16736	54676
	[276573]	[44127]	[89404]
Number of MSAs	193	193	193

Notes: Column 1 reports average metropolitan statistical area (MSA) characteristics in 1980. Column 2 reports the average change (in logs) in MSA characteristics from 1980 to 1990, and Column 3 reports the average change (in logs) from 1980 to 2000. Standard deviations reported in brackets.

Appendix Table 2. Estimated Impacts on Composition-Adjusted Earnings and Housing Costs

	Cross-section in 1980 and 1990 OLS (1)	Change from 1980 to 1990 OLS (2)	Change from 1980 to 2000 OLS (3)	Change from 1980 to 1990 2SLS (4)	Change from 1980 to 2000 2SLS (5)
Panel A. Log Earnings	0.28***	0.13*	0.20**	0.90***	1.15***
	(0.07)	(0.07)	(0.08)	(0.29)	(0.32)
Panel B. Log Cost of Rent	0.99***	0.28*	0.47***	1.96**	2.41***
	(0.18)	(0.16)	(0.17)	(0.77)	(0.71)
Panel C. Log Home Value	1.19***	0.52**	0.74***	1.84**	2.47
	(0.35)	(0.26)	(0.22)	(0.74)	(0.82)

Notes: In Panel A, each column reports estimates analogous to those in Panel A of Table 2, but adjusting annual earnings for worker composition by controlling for age, age squared, education (high school, some college, college), race, and gender. In Panels B and C, each column reports estimates analogous to those in Panels B and C of Table 2, but adjusting housing costs for physical characteristics by controlling for the number of rooms and number of bedrooms (dummy variables for each number), whether the home is part of a multi-unit structure, and the presence of a kitchen or plumbing. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix Table 3. Estimated Impacts of TFP, Alternative Specifications

	Outcome Change from 1980 to 2000:		
	TFP Change from 1980 to 1990 Control for TFP Change from 1990 to 2000	TFP Change from 1980 to 1990 IV for TFP Change from 1990 to 2000	TFP Change from 1980 to 2000
	2SLS	2SLS	2SLS
	(1)	(2)	(3)
Panel A. Log Earnings	1.31***	1.05*	0.75***
	(0.40)	(0.60)	(0.28)
Panel B. Log Cost of Rent	2.12***	2.11**	1.12***
	(0.59)	(0.89)	(0.41)
Panel C. Log Home Value	2.21***	2.47**	1.14**
	(0.70)	(1.17)	(0.47)
Panel D. Log Real Earnings	0.67**	0.42	0.41**
	(0.28)	(0.40)	(0.18)
Panel E. Log Workers	3.73***	4.37**	1.79***
	(1.09)	(1.80)	(0.61)

Notes: Column 1 reports estimates from equation 20 in the text (and Column 5 of Table 2), but controlling for the change in TFP from 1990 to 2000. Column 2 reports estimates from the same specification, but instrumenting for the change in TFP from 1990 to 2000 with the predicted change in TFP from 1990 to 2000. Column 3 reports estimates from a long-difference specification, regressing changes in each outcome on changes in TFP from 1980 to 2000, and instrumenting using the predicted change in TFP from 1980 to 2000. Each specification controls for region fixed effects and weights each MSA by its total manufacturing output in 1980. In Panel D, Log Real Earnings are defined as: Log(Earnings) - 0.3*Log(Cost of Rent). Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix Table 4. Estimated Impacts of TFP, Alternative Weightings

	Cross-section in 1980 and 1990 OLS (1)	Change from 1980 to 1990 OLS (2)	Change from 1980 to 2000 OLS (3)	Change from 1980 to 1990 2SLS (4)	Change from 1980 to 2000 2SLS (5)
Panel A. Population-Weighted Estimates					
Log Earnings	0.50***	0.12*	0.23*	0.91***	1.50***
	(0.07)	(0.07)	(0.12)	(0.33)	(0.46)
Log Cost of Rent	1.21***	0.30**	0.39***	1.62*	2.15***
	(0.15)	(0.15)	(0.14)	(0.87)	(0.69)
Log Home Value	1.75***	0.45*	0.59***	1.31	2.21***
	(0.26)	(0.24)	(0.21)	(1.05)	(0.80)
Log Real Earnings	0.14***	0.03	0.11	0.42**	0.86***
	(0.04)	(0.05)	(0.09)	(0.20)	(0.32)
Log Workers	5.95***	-0.03	-0.05	2.89***	5.06***
	(0.98)	(0.20)	(0.24)	(1.09)	(1.71)
Panel A. Unweighted Estimates					
Log Earnings	0.31***	0.02	0.03	0.56***	0.90***
	(0.06)	(0.06)	(0.07)	(0.22)	(0.28)
Log Cost of Rent	0.69***	0.06	0.06	1.12***	1.17***
	(0.11)	(0.09)	(0.08)	(0.39)	(0.36)
Log Home Value	0.74***	0.05	0.15	1.18***	1.36***
	(0.16)	(0.12)	(0.11)	(0.45)	(0.40)
Log Real Earnings	0.10**	-0.00	0.01	0.23*	0.55***
	(0.04)	(0.04)	(0.06)	(0.13)	(0.19)
Log Workers	1.88***	-0.01	0.13	1.97***	3.24***
	(0.54)	(0.16)	(0.20)	(0.68)	(0.93)

Notes: Panel A reproduces estimates from Table 2, but weighting each MSA by its total population in 1980. Panel B reproduces these estimates for unweighted specifications. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix Table 5. Robustness of the Interaction Between TFP and Elasticity of Housing Supply

	Cross-section in 1980 and 1990 OLS (1)	Change from 1980 to 1990 OLS (2)	Change from 1980 to 2000 OLS (3)	Change from 1980 to 1990 2SLS (4)	Change from 1980 to 2000 2SLS (5)
Panel A. All Baseline Specifications					
Log Earnings	-0.101*	-0.094***	-0.202***	-0.099**	-0.203***
	(0.057)	(0.035)	(0.052)	(0.048)	(0.070)
Log Cost of Rent	-0.368***	-0.330***	-0.291***	-0.372***	-0.298***
	(0.094)	(0.055)	(0.074)	(0.089)	(0.094)
Log Home Value	-0.620***	-0.456***	-0.416***	-0.478***	-0.400***
	(0.150)	(0.099)	(0.104)	(0.109)	(0.141)
Log Real Earnings	0.010	0.005	-0.115***	0.013	-0.114**
	(0.039)	(0.024)	(0.039)	(0.027)	(0.049)
Panel B. Controls for Region Interacted With TFP					
Log Earnings	-0.033	-0.080**	-0.178***	-0.093***	-0.193***
	(0.064)	(0.034)	(0.047)	(0.032)	(0.048)
Log Cost of Rent	-0.304***	-0.320***	-0.272***	-0.367***	-0.292***
	(0.108)	(0.062)	(0.084)	(0.061)	(0.080)
Log Home Value	-0.570***	-0.427***	-0.384***	-0.458***	-0.392***
	(0.152)	(0.100)	(0.100)	(0.102)	(0.103)
Log Real Earnings	0.059	0.016	-0.097***	0.018	-0.106***
	(0.042)	(0.022)	(0.033)	(0.021)	(0.034)
Panel C. All Baseline Specifications, Omitting MSAs with 10 Highest and 10 Lowest Housing Elasticities					
Log Earnings	-0.013	-0.026	-0.142***	-0.107	-0.184**
	(0.074)	(0.038)	(0.054)	(0.066)	(0.075)
Log Cost of Rent	-0.231**	-0.255***	-0.229**	-0.461***	-0.290***
	(0.115)	(0.093)	(0.108)	(0.144)	(0.104)
Log Home Value	-0.193	-0.253*	-0.314***	-0.381**	-0.363**
	(0.134)	(0.135)	(0.113)	(0.173)	(0.154)
Log Real Earnings	0.056	0.051**	-0.074**	0.032	-0.097*
	(0.059)	(0.023)	(0.034)	(0.035)	(0.053)
Panel D. All Baseline Specifications, Alternative Measure of Housing Elasticity from Gyourko, Saiz, and Summers (2008)					
Log Earnings	-0.137***	-0.165***	-0.277***	-0.282***	-0.429***
	(0.050)	(0.033)	(0.051)	(0.064)	(0.092)
Log Cost of Rent	-0.369***	-0.286***	-0.196***	-0.510***	-0.414***
	(0.098)	(0.065)	(0.072)	(0.121)	(0.119)
Log Home Value	-0.716***	-0.524***	-0.420***	-0.719***	-0.629***
	(0.159)	(0.094)	(0.101)	(0.153)	(0.157)
Log Real Earnings	-0.027	-0.079***	-0.218***	-0.129***	-0.305***
	(0.040)	(0.023)	(0.037)	(0.037)	(0.065)

Notes: Panel A reports estimates that correspond to the interaction effects in Table 3, but for both OLS and 2SLS specifications. Panel B includes controls for region interacted with TFP (or changes in TFP), which focuses on within-region variation in housing elasticity. Panel C omits MSAs with the 10 highest and 10 lowest measures of the elasticity of housing supply. Panel D uses an alternative measure of housing elasticity from the Wharton data, which is also normalized to have a standard deviation of one. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix Table 6. OLS Estimated Impacts of TFP, by Education Level

	Cross-Section (OLS):			Short-run Effect: Change from 1980 to 1990 (OLS)			Long-run Effect: Change from 1980 to 2000 (OLS)		
	College	Some College	No College	College	Some College	No College	College	Some College	No College
	(1)	(2)	(3)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. Log Earnings	0.28***	0.25***	0.11**	0.13**	0.14**	0.15**	0.23**	0.21***	0.18**
	(0.06)	(0.06)	(0.06)	(0.05)	(0.06)	(0.06)	(0.09)	(0.08)	(0.08)
Panel B. Log Cost of Rent	0.79***	0.78***	0.80***	0.18	0.34**	0.32**	0.35**	0.39***	0.45***
	(0.17)	(0.14)	(0.13)	(0.15)	(0.14)	(0.16)	(0.16)	(0.13)	(0.14)
Panel C. Log Home Value	0.88***	0.95***	1.04***	0.47**	0.53**	0.52**	0.63***	0.66***	0.71***
	(0.22)	(0.23)	(0.24)	(0.19)	(0.23)	(0.25)	(0.17)	(0.18)	(0.20)
Panel D. Log Real Earnings	0.04	0.01	-0.13**	0.07	0.04	0.05	0.13*	0.10*	0.04
	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)	(0.05)	(0.07)	(0.06)	(0.06)
Panel E. Log Workers	4.60***	3.80***	3.18***	-0.05	0.03	0.11	0.26	0.16	0.06
	(1.15)	(1.05)	(1.06)	(0.21)	(0.18)	(0.18)	(0.27)	(0.29)	(0.25)

Notes: The reported estimates are analogous to those in Table 4, but correspond to the OLS specifications (from Table 2) that were not shown in Table 4, which was restricted to our preferred 2SLS specifications. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix Table 7. OLS Estimated Impacts of TFP, by Sector

	Cross-Section (OLS):		Short-run Effect: Change from 1980 to 1990 (OLS)		Long-run Effect: Change from 1980 to 2000 (OLS)	
	Manufacturing	Non-Manufacturing	Manufacturing	Non-Manufacturing	Manufacturing	Non-Manufacturing
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Log Earnings	0.29***	0.38***	0.15*	0.10	0.30**	0.24**
	(0.09)	(0.09)	(0.08)	(0.07)	(0.14)	(0.11)
Panel B. Log Workers	3.10***	3.79***	0.09	0.05	0.22	0.18
	(0.94)	(1.16)	(0.21)	(0.16)	(0.28)	(0.22)

Notes: The reported estimates are analogous to those in Table 6, but correspond to the OLS specifications (from Table 2) that were not shown in Table 5, which was restricted to our preferred 2SLS specifications. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.