

ESTIMATING WHO BENEFITS FROM PRODUCTIVITY GROWTH: LOCAL AND DISTANT EFFECTS OF CITY PRODUCTIVITY GROWTH ON WAGES, RENTS, AND INEQUALITY

Richard Hornbeck and Enrico Moretti*

Abstract—We first estimate the direct effects on local workers' earnings and housing costs from increases in local labor demand caused by gains in city-level manufacturing productivity. We find that local workers benefit from productivity growth, even after subtracting increases in housing costs. These gains are larger for local less educated workers, such that productivity growth reduces local inequality. We then propose and implement a new transparent method of estimating indirect effects of local productivity growth on earnings and housing costs of workers in other cities. We find that these general equilibrium effects are economically important and disproportionately benefit college-educated workers.

I. Introduction

ECONOMISTS have long considered differences in productivity across countries to be a key source of differences in standard of living. Indeed, increases in consumption and welfare depend centrally on productivity growth in many macroeconomic models (see Solow, 1956), although empirical evidence is mixed on the causal link between productivity and real earnings.¹ There is tremendous variation in productivity growth across cities within the United States, and workers' migration responses to these productivity differences can change the distribution and magnitude of gains from productivity growth.

In this paper, we examine differences in revenue total factor productivity growth (TFPR) across U.S. metropolitan areas (MSAs) and estimate the magnitude and distribution of effects on workers' earnings and housing costs.

Received for publication July 16, 2020. Revision accepted for publication January 18, 2022. Editor: Amit K. Khandelwal.

*Hornbeck: University of Chicago and NBER; Moretti: University of California, Berkeley, and NBER.

For helpful comments and suggestions, we thank many colleagues and seminar participants at BEA, Bocconi, Chicago Economics, Chicago Harris, CUHK, European Central Bank, LSE, NBER (CRIW, Economic Growth, Labor, Urban), NYU, Oxford, Paris School of Economics, Stanford, University of Bologna, USC, and Yale. Andrea Cerrato, William Cockriel, Julius Luetge, and Joseph Root provided extended research assistance, with additional assistance by Georgios Angelis, Melissa Eccleston, Alonso Sanchez, and Alex Weckenman. Any opinions and conclusions expressed here are our own and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01208.

¹In country-level data, the correlation between wage growth and productivity growth appears to have weakened (Jones & Klenow, 2016; Schwelonus, Kappeler, & Pionnier, 2017; Stansbury & Summers, 2017).

We explore how within-country productivity differences influence workers' standard of living and the division of economic gains between labor and land, in contrast to more recent focus on the division of economic gains between labor and capital. On a methodological level, we propose a new approach to estimate general equilibrium effects of local shocks.

The incidence of local productivity growth, or who benefits to what degree, is not clear ex ante. City TFPR growth raises local labor demand and therefore is likely to increase earnings for local workers. When city TFPR growth induces substantial in-migration from other cities, however, this may raise housing costs and limit wage gains for local workers who rent their homes. Because housing and other nontradable goods account for the majority of worker consumption, changes in these local costs have potential large consequences for workers' standard of living (Moretti, 2013; Diamond & Moretti, 2022). Substantial increases in housing costs in increasingly productive cities revive the classical concern of Ricardo, in which landowners capture the gains from productivity growth.

We measure city-level changes in TFPR using data from the U.S. Census of Manufactures. The manufacturing sector experienced gains in TFPR from 1980 to 1990, and we uncover substantial geographic differences in TFPR growth across cities. We use this cross-city heterogeneity in TFPR growth, from 1980 to 1990, to estimate its effects on city employment, earnings, and housing costs using U.S. Census data from 1980 to 2010. We focus on the manufacturing sector because in the period we focus on, the manufacturing sector was at its peak and employed 20 million Americans (Charles, Hurst, & Notowidigdo, 2018), accounting for the majority of employment in the tradable sector.

An important aspect of our analysis is that we do not limit our focus to direct effects of local TFPR growth on earnings and housing costs in that same city. We propose and implement a novel way to estimate indirect effects that local productivity growth has on other cities through worker mobility. Empirically, these general equilibrium effects are quantitatively important for characterizing the magnitude and distribution of who benefits from local productivity growth.

In the first part of the paper, we estimate direct effects of TFPR growth on cities experiencing relatively larger TFPR

growth. Because of the potential for local shocks to create spurious associations between TFPR growth and wages or housing costs, our empirical approach uses four instrumental variables to predict changes in city-level TFPR. Our baseline instrumental variable reflects industry-specific changes in nationwide TFPR that have differential effects on cities due to differences in cities' initial industry concentration. We construct an alternative instrument, based on changes in stock prices by industry, to proxy for unexpected shocks to industry TFPR that may differentially affect cities concentrated in those industries. A third instrument is based on changes in exposure to export markets, since trade exposure affects firm output prices and has been associated with patenting and investment in R&D (Aghion et al., 2020; Autor et al., 2020). A fourth instrument is based on patenting activity within technology classes, which disproportionately increases TFPR in cities that had been patenting more in those technology classes.

The instruments have a similar structure, but the underlying identification assumptions are instrument specific. Importantly, the instruments use different sources of empirical variation: the cities that are predicted to have larger TFPR changes for one of our instruments are different from the cities predicted to have larger TFPR changes for other instruments. In addition, the most influential industries for predicting MSA-level TFPR growth are often different across the instruments (Goldsmith-Pinkham, Sorkin, & Swift, 2020). The alternative instruments yield similar estimates, however, and overidentification tests fail to reject that the estimates are statistically indistinguishable.

We estimate that local TFPR growth increases the earnings of local workers. A 1% increase in city-level manufacturing TFPR from 1980 to 1990 is associated with an average long-run increase of 1.5% in annual earnings from 1980 to 2000. Local employment increases by 4%, driven by in-migration. As a consequence of this in-migration, demand for housing increases. We estimate that a 1% increase in city-level manufacturing TFPR is associated with a 1.5% increase in housing rents and a 2.5% increase in home values.

Who benefits from local TFPR growth then depends in large part on residents' position in the housing market. For workers who rent their home, much of the increase in earnings is offset by increases in the local cost of living. For workers who had owned their home, the gains are much larger because they come through both higher wages and higher housing values. We calculate impacts on worker purchasing power, which reflects earnings adjusted for cost of living, and find that a 1% increase in local TFPR increases purchasing power of renters and homeowners by 0.6% and 1.6%, respectively.

Who benefits locally from TFPR growth also depends substantially on workers' level of education. We estimate greater impacts on both nominal earnings and purchasing power for high school graduates than for college graduates. At the same time, increases in local employment are substantially larger for more educated workers. We interpret the larger earnings gains and smaller employment effects experi-

enced by high school graduates as caused by their lower geographical mobility compared to college graduates (Bound & Holzer, 2000; Wozniak, 2010; Malamud & Wozniak, 2012; Notowidigdo, 2020). Since less educated workers are less geographically mobile, their local supply is more inelastic, and so the incidence of local TFPR growth falls more on less educated workers. An important implication is that local TFPR growth compresses local inequality. We estimate that increases in nominal earnings and purchasing power are substantially greater for workers at the 10th percentile and 50th percentile of the income distribution than for workers at the 90th percentile of the income distribution.

In the second part of the paper, we turn to the indirect effects of local TFPR growth on earnings and housing costs in other cities. To see why indirect effects may be important, consider an example in which Houston experiences a positive TFPR shock that raises wages and employment in Houston. As migrants to Houston leave other places, this puts upward pressure on wages and downward pressure on housing costs in those other places (given a downward-sloping labor demand and upward-sloping housing supply). These indirect effects are diffused when a small share of migrants to Houston come from each other place, but the sum of these indirect effects may be quantitatively important alongside the direct effect in Houston.

To estimate indirect effects from worker mobility, for each city hit by a TFPR shock, we use our estimates of the direct employment effects and assumptions on city-to-city migrant flows to estimate employment declines in other cities due to out-migration. We then use data on city-level elasticities of housing supply, along with an assumption on the elasticity of labor demand, to gauge the general magnitude of these indirect effects on housing costs and earnings in other cities.

The typical approach in the literature to estimate general equilibrium effects of localized shocks is to impose a quantitative spatial model, identifying general equilibrium effects by drawing on the structure of the model. An advantage of our approach is that it requires fewer assumptions. A disadvantage is that we do not explore impacts on workers through changes in traded goods prices or impacts on firms and shareholders, which are components of the aggregate impacts of TFPR growth and depend on how much TFPR growth is driven by physical productivity improvements or increases in firm pricing power.

We find indirect effects that are economically substantial and have important implications for understanding who ultimately benefits from local TFPR growth. We estimate that 38% of the combined increase in purchasing power for the average worker occurs outside cities directly affected by local TFPR growth. These indirect effects disproportionately benefit college-educated workers due to their greater geographic mobility, which counterbalances the local decrease in inequality associated with local TFPR growth. Further, these indirect effects disproportionately benefit renters in other cities, who benefit from decreased housing costs.

Overall, when we sum the direct effects and indirect effects from local TFPR growth, we find that the average

worker benefited substantially from manufacturing TFPR growth that averaged 5.3% from 1980 to 1990. We estimate that this increase in real TFPR led to economically large increases in purchasing power: on the order of 0.5% to 0.6% per year, between 1980 and 2000, for the average full-time worker.

The summed direct effects and indirect effects are roughly similar for more educated workers and less educated workers, in percentage terms. Less educated workers benefit more from TFPR growth in their city, but more educated workers benefit more from TFPR growth in other cities. Thus, neglecting indirect effects from worker mobility would both understate the gains from local TFPR growth and also yield a biased view of its distributional consequences. These estimates complement the large literature on skill-biased technological change and labor-saving technological change, which explores increases in inequality from productivity growth.²

The gains from TFPR growth are very different, however, depending on workers' geographical location. The benefits of productivity growth are economically large in cities that benefit from strong direct and indirect effects (e.g., San Jose) and minimal in cities with weak direct and indirect effects (e.g., St. Louis). Thus, on net, the average worker benefits substantially from productivity growth, but these gains depend in large part on where the worker lives. A high-level view of average country-level changes would mask substantial variation in experiences across areas and people.

From a methodological point of view, our results suggest caution for interpreting empirical results that focus exclusively on the local impacts of local shocks. Many studies in economics seek to estimate the effects of economic shocks, such as immigration (Card, 2001) or trade (Autor, Dorn, & Hanson, 2013), by comparing areas that experience large shocks to areas that do not. Our findings indicate that when local shocks generate large migration responses, a substantial portion of the overall effects may be missed when focusing only on the direct effects. Including these indirect effects, even those indirect effects from worker mobility only, can yield qualitatively and quantitatively different conclusions. Our approach, based on migrant flows and certain key elasticities, can be used in other contexts to gauge the magnitude of indirect effects in a reduced-form manner. General equilibrium models can quantify additional indirect effects, such as impacts on prices of traded goods, though at the cost of imposing stronger theoretical assumptions on the structure of the economy.

II. Data

For 193 metropolitan statistical areas (MSAs), covering 63% of the U.S. population in 1980, we combine data from the Census of Population and the Census of Manufactures.

²In contrast to skill-biased technological change that favors more skilled workers or labor-saving technological change that potentially reduces labor demand, TFPR growth is skill neutral and raises labor demand.

We measure labor market outcomes and housing market outcomes at the MSA level, aggregating individual-level data and household-level data from the 1980, 1990, and 2000 waves of the Census of Population and the five-year sample from the American Community Survey centered on 2010 (2008–2012). The main outcome variables are average annual earnings, average household gross rent (for renters), average household home value (for owners), and city employment. For all outcomes, we analyze city-level averages and separate city-level averages within education group.³ Appendix table 1 reports average characteristics of the 193 sample cities in 1980 and average changes over time.

We measure average city-level TFPR using confidential plant-level data from the Census of Manufactures (CMF) in 1977, 1987, and 1997. The CMF contains plant-level data on all manufacturing plants' employment, capital stock, materials, and total value of shipments. We refer to years 1980, 1990, and 2000 with the understanding that these data are measured three years prior.


To estimate average city-level TFPR, in each decade, we adopt an econometric approach that is similar to that used in our previous work based on the same data from the Census of Manufactures (Greenstone, Hornbeck, & Moretti, 2010). We assume each plant uses a Cobb-Douglas technology and, in each year separately, we regress log output on log input expenditures and city fixed effects (weighting by plant output). The estimated 193 city fixed effects, in each decade, reflect average TFPR in that city and decade. In appendix A, we report details of the estimation procedure and its limitations.

Our measure of TFPR is a measure of "revenue productivity," as is typical in the literature, and therefore productivity growth in our context reflects increased value of plant output given plant input expenditures. This reflects not only physical productivity increases (more quantity produced for a given set of inputs), but also relative increases in output prices (e.g., due to increased demand for firm output). This measure of revenue productivity captures changes in local labor demand, the main focus of our paper, and both sources of variation in TFPR (from prices or physical productivity) have equivalent effects on local labor markets and local housing markets through increases in firm labor demand. Consumers of traded goods are affected differently by increases in physical productivity or increased firm prices, as are firm owners, but we do not analyze these other consequences of productivity growth.


³We use a sample of adult full-time workers. Following standard practice (see Katz & Autor, 1999), we restrict 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 at least 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. We multiply top-coded earnings and home values by 1.5, which make up 0.26% of observations for earnings in 1980 and 0.69% to 2.87% of observations for home values from 1980 to 2010. The estimates are nearly identical without this adjustment or when multiplying by 2 instead of 1.5. Other reported values reflect state-level mean or median values above the Census-reporting threshold.

FIGURE 1.—SPATIAL DISTRIBUTION OF REVENUE TOTAL FACTOR PRODUCTIVITY, 1980 AND 1990


Panel A. TFPR in 1980



Panel B. TFPR in 1990



Panel C. Change in TFPR from 1980 to 1990



Panel D. Total Manufacturing Output by MSA



Panels A and B show revenue total factor productivity (TFPR) in 1980 and 1990 for the 193 sample MSAs, and panel C shows the change in TFPR from 1980 to 1990. MSAs are separated into 10 groups, with darker-shaded groups representing MSAs with greater TFPR (or a greater relative change in TFPR). Panel D shows manufacturing output for each sample MSA in 1980, with darker shades representing greater manufacturing output.

For each city, figure 1 shows average TFPR in 1980 (panel A), in 1990 (panel B), and in changes from 1980 to 1990 (panel C). There is substantial variation in TFPR growth across cities, within broader geographic regions, that we use in the empirical analysis.⁴ Reassuringly, though, there is also persistence in TFPR across areas, with higher productivity places in 1980 remaining higher productivity in 1990 (appendix figure 1). Appendix figure 2A shows that measured local TFPR growth from 1980 to 1990 is largely uncorrelated with measured local TFPR growth from 1990 to 2000. Panels B, C, and D show that local TFPR growth from 1980 to 1990 is weakly positively correlated with local TFPR growth in cities within 100 miles and not correlated with local TFPR growth in cities within 250 miles and 500 miles.

III. Empirical Specifications and Identification of Direct Effects

To estimate the direct effects of local TFPR growth, we relate changes in city TFPR to changes in that city's labor and housing market outcomes (employment, earnings, hous-

ing costs). Local TFPR growth increases local labor demand, which increases nominal wages and the cost of housing. The local gains from TFPR growth are then split between workers and landowners: the incidence depends on relative elasticities and which of the two factors (labor or housing) is supplied more elastically. Appendix B presents a simple spatial equilibrium model (Rosen-Roback) that helps interpret how TFPR growth in a city may affect employment, wages, and housing costs in that city and indirectly affect other cities through worker mobility.

We regress the change in outcome Y_c in city c (employment, earnings, housing costs) on the change in city TFPR A_c :

$$\begin{aligned} & \ln(Y_{c,1990}) - \ln(Y_{c,1980}) \\ &= \pi^M (\ln \hat{A}_{c,1990} - \ln \hat{A}_{c,1980}) + \alpha_r + \varepsilon_c, \end{aligned} \quad (1)$$

$$\begin{aligned} & \ln(Y_{c,2000}) - \ln(Y_{c,1980}) \\ &= \pi^L (\ln \hat{A}_{c,1990} - \ln \hat{A}_{c,1980}) + \alpha_r + \varepsilon_c, \end{aligned} \quad (2)$$

$$\begin{aligned} & \ln(Y_{c,2010}) - \ln(Y_{c,1980}) \\ &= \pi^{XL} (\ln \hat{A}_{c,1990} - \ln \hat{A}_{c,1980}) + \alpha_r + \varepsilon_c, \end{aligned} \quad (3)$$

where \hat{A}_{ct} is our estimate of average TFPR in city c in year t and α_r are Census region fixed effects. We consider three

⁴The changes in TFPR at different parts of the distribution are -2.2% (10th percentile), 0.4% (25th percentile), 4.7% (50th percentile), 10.8% (75th percentile), and 13.7% (90th percentile).

time horizons: medium run (change in outcomes between 1980 and 1990), long run (change in outcomes between 1980 and 2000), and longer run (change in outcomes between 1980 and 2010). In all three cases, TFPR growth is from 1980 to 1990. Thus, these specifications allow for additional time in reaching a new spatial equilibrium, as workers and firms relocate and there is construction of new housing units. In addition, in the presence of agglomeration spillovers, the effects of local TFPR growth may increase over time due to self-reinforcing dynamics. Across all specifications, we report robust standard errors adjusted for heteroskedasticity.⁵

These equations reflect reduced-form relationships between city TFPR and city-level outcomes, which we expect to depend on the relative elasticities of local labor supply and housing supply (appendix B). We explore how the estimated impacts of local TFPR growth vary with worker education, as more educated workers are generally more geographically mobile and the more elastic supply of more educated workers may reduce the impact on local earnings of more educated workers and compress local inequality. We also explore how the impacts vary with cities' elasticity of housing supply (Saiz, 2010).

OLS estimation of equations (1 to 3) is likely to be biased for two categories of reasons. First, estimated changes in TFPR are likely to contain substantial measurement error. TFPR is a residual, measured with error, and the empirical specifications examine changes in TFPR that exacerbate bias from measurement error. Second, changes in city-level TFPR may be influenced by changes in local factors that independently affect employment, local earnings, or housing costs. These biases could be either positive or negative, given changes in productive amenities or consumption amenities.⁶

A. Instrumental Variables

We instrument for changes in city-level TFPR using four alternative instrumental variables, in isolation or in combination, with the goal of isolating changes in local TFPR that reflect national influences that disproportionately affect cities based on their initial characteristics.

Our baseline instrumental variable uses nationwide changes in TFPR by industry to predict each city's change

⁵The regressions are weighted by each city's total manufacturing output in 1980. Our measure of TFPR reflects data grouped at the city 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. The pattern of estimates is similar for unweighted regressions and somewhat smaller in magnitude, consistent with smaller effects on smaller MSAs (appendix table 9).

⁶For example, an improvement in local transportation infrastructure could both increase local TFPR and the desirability of the area for workers, which would cause OLS estimates to understate the impact of TFPR on wages and overstate the impact of TFPR on rents. On the other hand, tighter air quality regulations may lower TFPR, decrease nominal wages, and increase housing costs, causing OLS estimates to overstate the impact of TFPR on wages and understate the impact of TFPR on rents.

in TFPR depending on each city's initial concentration of industries. For each city, the instrument is defined by summing over all three-digit SIC industries: the city's 1980 fraction of manufacturing output in an industry ($\alpha_{i,c,1980}$), multiplied by the national change in TFPR for that industry from 1980 to 1990 ($\gamma_{i,c,1980-1990}$), such that $IV_c^{\text{baseline}} = \sum_i \alpha_{i,c,1980} \times \gamma_{i,c,1980-1990}$. The national change in TFPR for industry i is indexed by city c because, to avoid mechanical correlation between industry-level changes and city-specific shocks, we omit that particular city and estimate "leave-out" national changes in TFPR by industry across all other cities. For each city, the predicted change in TFPR from 1980 to 1990 then depends on that city's industries in 1980 and changes in TFPR from 1980 to 1990 for those industries in other parts of the country.

A second instrument uses industry-level stock market returns to capture a variety of factors, including improvements in production technologies and increased demand for firm output, which are associated with increased revenue productivity of particular industries and increased labor demand. These industry-level gains may then benefit most those cities that were initially concentrated in those industries: we calculate industry-specific stock market returns from 1980 to 1990 ($\gamma_{i,c,1980-1990}^s$), assigned to cities based on their industry employment shares in 1980 ($\alpha_{i,c,1980}^s$), such that $IV_c^{\text{stock}} = \sum_i \alpha_{i,c,1980}^s \times \gamma_{i,c,1980-1990}^s$.

One notable feature of this instrument, which may also be useful when applied to other empirical contexts, is that relative changes in stock prices between 1980 and 1990 are arguably unpredictable in 1980. This is in contrast to the baseline instrument and other instruments, which use variation that may be partially predicted at the beginning of the period. Thus, a comparison of estimates based on this instrument with estimates based on the other instruments is informative about how much the estimates based on the other instruments reflect unexpected changes.

The third instrument is based on increased industry exposure to export markets, which may increase TFPR for two reasons. First, increased export demand may translate into higher output prices and therefore higher revenue productivity in cities initially more concentrated in those industries. Second, increased net imports have been found to reduce innovation in U.S. firms (Autor et al., 2020); conversely, increased net exports may have a positive effect (Aghion et al., 2020). We aim to isolate exogenous trade shocks to U.S. industries by measuring increases in exports from other high-income countries. This instrument is calculated as the product of baseline city industry employment shares ($\alpha_{i,c,1980}^e$) times the change in exports by industry from 1980 to 1990 ($\gamma_{i,1980-1990}^e$), such that $IV_c^{\text{export}} = \sum_i \alpha_{i,c,1980}^e \times \gamma_{i,1980-1990}^e$. The instrument then reflects a city-specific index of export

⁷We calculate an index of stock market returns by industry from 1980 to 1990 using monthly CRSP data. When assigning industry-specific growth rates to a city, we exclude companies headquartered in that city.

exposure, based on a weighted average of industry-specific growth in exports per worker.⁸

A fourth instrument is based on patenting activity. Cities initially concentrated in particular technology classes may experience greater TFPR growth when there is greater patenting activity in those technologies nationwide.⁹ For each city, the instrument is defined by summing over all technology classes i : the number of patent assignees per manufacturing worker in 1980 ($\alpha_{i,c,1980}^p$, multiplied by the total number of patents filed nationwide between 1980 and 1990 ($\gamma_{i,c,1980-1990}^p$) excluding patents from an assignee located in that city, such that $IV_c^{\text{patent}} = \sum_i \alpha_{i,c,1980}^p \times \gamma_{i,c,1980-1990}^p$. This instrument captures the relative benefit of national patenting activity for cities that had more patenting activity in particular technology classes.¹⁰

The identification assumption is specific to each instrument, though the four instruments are constructed similarly, as the assumptions depend on which cities are disproportionately affected by which national changes. For our baseline instrument, the identification assumption is that changes in labor market and housing market outcomes in certain cities (with manufacturing output initially concentrated in industries that experience stronger nationwide TFPR gains) would otherwise have been similar, on average, to changes in other cities (with manufacturing output initially concentrated in industries that experience weaker nationwide TFPR gains). For the alternative instruments, the identification assumptions are that labor market and housing market outcomes would otherwise have changed similarly in cities that were differentially exposed to stock market appreciation, export growth, or patenting activity.

These identification assumptions are meaningfully distinct when, in practice, the four instrumental variables capture different sources of empirical variation in city TFPR growth. Table 1, panels A and B, shows the sample cities with the largest and smallest predicted changes in TFPR for each of the four instrumental variables. While there is some overlap in these lists, the cities predicted to experience the

greatest TFPR growth between 1980 and 1990 based on the baseline instrument are not the same set of cities predicted to experience the greatest TFPR growth based on the stock market instrument, export instrument, or patent instrument. For example, the top three cities in predicted TFPR growth are all different across the baseline instrument (Richmond, Atlantic City, and Raleigh-Durham), stock market instrument (Greenville, Charlotte, Greensboro), export instrument (Lexington, Fort Collins, Binghamton), and patent instrument (Stamford, Washington, Wilmington). There is more overlap among cities predicted to experience the least TFPR growth between 1980 and 1990, particularly among cities with substantial exposure to the oil and gas industry that experienced negative shocks in the 1980s, and we later estimate that the estimates are not sensitive to controlling for cities' baseline share in the oil and gas industry.

Table 1, panel C, lists which industries or technology classes are most influential in driving the variation for each instrument.¹¹ The instruments largely reflect nationwide shocks to different industries, which then affect different cities according to their baseline concentrations. Petroleum refining is influential for both the baseline IV and export IV, which may reflect shocks in the oil and gas industry that affect city TFPR through multiple channels, but our identification assumption does not rely on city TFPR changing only through differential exposure to trade shocks or other shocks.

Figure 2 shows pairwise correlations for all pairs of the four instruments, where each dot represents a city. The different instruments are statistically correlated in three of the six cases but reflect a great amount of independent variation, with pairwise regressions yielding R -squared values of 0.002, 0.358, 0.123, 0.014, 0.006, and 0.006.

Figure 3 shows four maps, one for each instrument, that illustrate the geographic variation in predicted changes in TFPR. Each instrument is divided into deciles, with darker shades reflecting higher values of the instrument (which predict greater increases in TFPR). The maps show that the instruments are not simply picking up local shocks common to each instrument and that there is geographic variation within nearby areas for each instrument.

Because the instrumental variables reflect different sources of empirical variation in predicted TFPR growth, the identification assumptions are also then meaningfully distinct in practice. We estimate overidentification tests that fail to reject that the 2SLS estimates are statistically the same across these instruments. Each instrumental variable estimates a particular local average treatment effect (LATE),

⁸Export data are from the UN Comtrade database (United Nations, 2003), which include industry exports from 28 high-income countries (excluding the United States) to 94 countries of all income levels. We calculate the growth in industry exports per worker, using the total number of workers in that industry across all cities in the United States in 1980. The weights are industry employment shares in each city from the 1980 Census. In the trade data, industry definitions are based on SITC Rev. 1 four-digit industries.

⁹The patent data are organized by technology class, and different technology classes experienced different rates of patenting over our period of analysis. For example, from 1980 to 1990, the three technology classes with the greatest patent assignees were Drug, Bio-Affecting and Body Treating Compositions, Stock Material or Miscellaneous Articles, and Measuring and Testing. We use patent data by technology class from the NBER Patent Data Project (Hall, Jaffe, & Trajtenberg, 2001). We match assignee location names to cities using the geographical correspondence engine of the Missouri Census Data Center.

¹⁰Of the four instruments, this patent instrument is the weakest in predicting city-level TFPR growth, perhaps because of skewed patent counts by technology class and city. We also give less emphasis to this instrument due to concerns that cities with greater patenting may otherwise have changed differently over this period.

¹¹For the Stock IV, Export IV, and Patent IV, we report the industries or technology classes with the highest Rotemberg weight in absolute value (following Goldsmith-Pinkham, Sorkin, & Swift, 2020). For the Baseline IV, as we no longer have access to the confidential plant-level data, we report which industry employment shares are most predictive of cities' predicted TFPR growth for the baseline instrument (i.e., the industry shares with the highest R -squared when regressing predicted TFPR growth on each industry share individually).

TABLE 1.—VARIATION IN PREDICTED CHANGES IN CITY TFPR, BY INSTRUMENTAL VARIABLE

	Baseline IV (1)	Stock Market IV (2)	Export IV (3)	Patent IV (4)
A. Cities with Greatest Predicted TFPR Growth				
1.	Richmond, VA	Greenville, SC	Lexington, KY	Stamford, CT
2.	Atlantic City, NJ	Charlotte, NC	Fort Collins, CO	Washington, DC
3.	Raleigh-Durham, NC	Greensboro, NC	Binghamton, NY	Wilmington, DE
4.	Little Rock, AR	Augusta, GA	Rochester, NY	Kalamazoo, MI
5.	Greeley, CO	Fayetteville, NC	Stamford, CT	Saginaw, MI
6.	Columbia, MO	Vineland, NJ	San Jose, CA	Albany, NY
7.	Lubbock, TX	El Paso, TX	Raleigh-Durham, NC	New Haven, CT
8.	Greensboro, NC	New Bedford, MA	Austin, TX	Trenton, NJ
9.	Pensacola, FL	Anniston, AL	Boise City, ID	New York, NY
10.	Austin, TX	McAllen, TX	Phoenix, AZ	Pittsburgh, PA
B. Cities with Least Predicted TFPR Growth				
1.	Bakersfield, CA	Eugene-Springfield, OR	Beaumont-Port Arthur, TX	Billings, MT
2.	Beaumont-Port Arthur, TX	Waterloo-Cedar Falls, IA	Corpus Christi, TX	Montgomery, AL
3.	Corpus Christi, TX	Detroit, MI	Billings, MT	Mobile, AL
4.	Billings, MT	Peoria, IL	Bakersfield, CA	Alexandria, LA
5.	Galveston-Texas City, TX	Odessa, TX	Galveston-Texas City, TX	Augusta-Aiken, GA-SC
6.	Baton Rouge, LA	Mobile, AL	Lafayette, LA	Abilene, TX
7.	Wichita, KS	Rockford, IL	Baton Rouge, LA	Nashville, TN
8.	Houston-Brazoria, TX	Davenport, IA	Houston-Brazoria, TX	Fayetteville, AR
9.	Lima, OH	Jackson, MI	Odessa, TX	McAllen-Edinburg, TX
10.	Odessa, TX	Beaumont-Port Arthur, TX	Anchorage, AK	Galveston-Texas City, TX
C. Most Influential Industries or Technology Classes for Each Instrument				
1.	Petroleum refining	Textile mill products (yarn, thread, carpets, rugs)	Petroleum refining	Stock material or miscellaneous articles
2.	Industrial and miscellaneous chemicals	Transportation equipment	Aircraft and parts	Synthetic resins or natural rubbers
3.	Tobacco manufactures	Metal products	Other primary metal industries	Adhesive bonding and misc. chemical manufacture
4.	Iron and steel foundries	Food products	Computers and related equipment	Metal working
5.	Motor vehicles and motor vehicle equipment	Printing and publishing	Miscellaneous fabricated metal products	Plastic and nonmetallic article shaping or treating processes

Panels A and B report the sample cities (MSAs) with the largest and smallest predicted growth in TFPR from 1980 to 1990 for each of the instrumental variables: baseline instrument (column 1), stock market instrument (column 2), export exposure instrument (column 3), and patent instrument (column 4). Panel C reports the industries or technologies with the highest estimated Rotemberg weight in absolute value (columns 2, 3, 4). For the baseline IV, in column 1, we report which industry shares have the highest R -squared when regressing predict TFPR growth on each industry share individually.

which reflects variation in TFPR growth due to sectoral shifts that may be more long-lasting than all observed variation in TFPR that is a combination of permanent shocks and transitory shocks. The IV estimates may then be larger than the OLS estimates, even in the absence of omitted variable bias or measurement error, though in practice, we see little systematic mean reversion in TFPR growth from one decade to the next (appendix figure 2). Our estimates may be expected to reflect the impacts of average TFPR growth, sustained over longer periods, whereas alternative research designs based on temporary TFPR shocks may identify different relationships (e.g., if there is little time for migration responses).

IV. Direct Effects of Local TFPR Growth

A. Direct Effects on Employment, Earnings, and Housing Costs

Table 2 reports our baseline estimates, which instrument for changes in TFPR using the baseline instrumental vari-

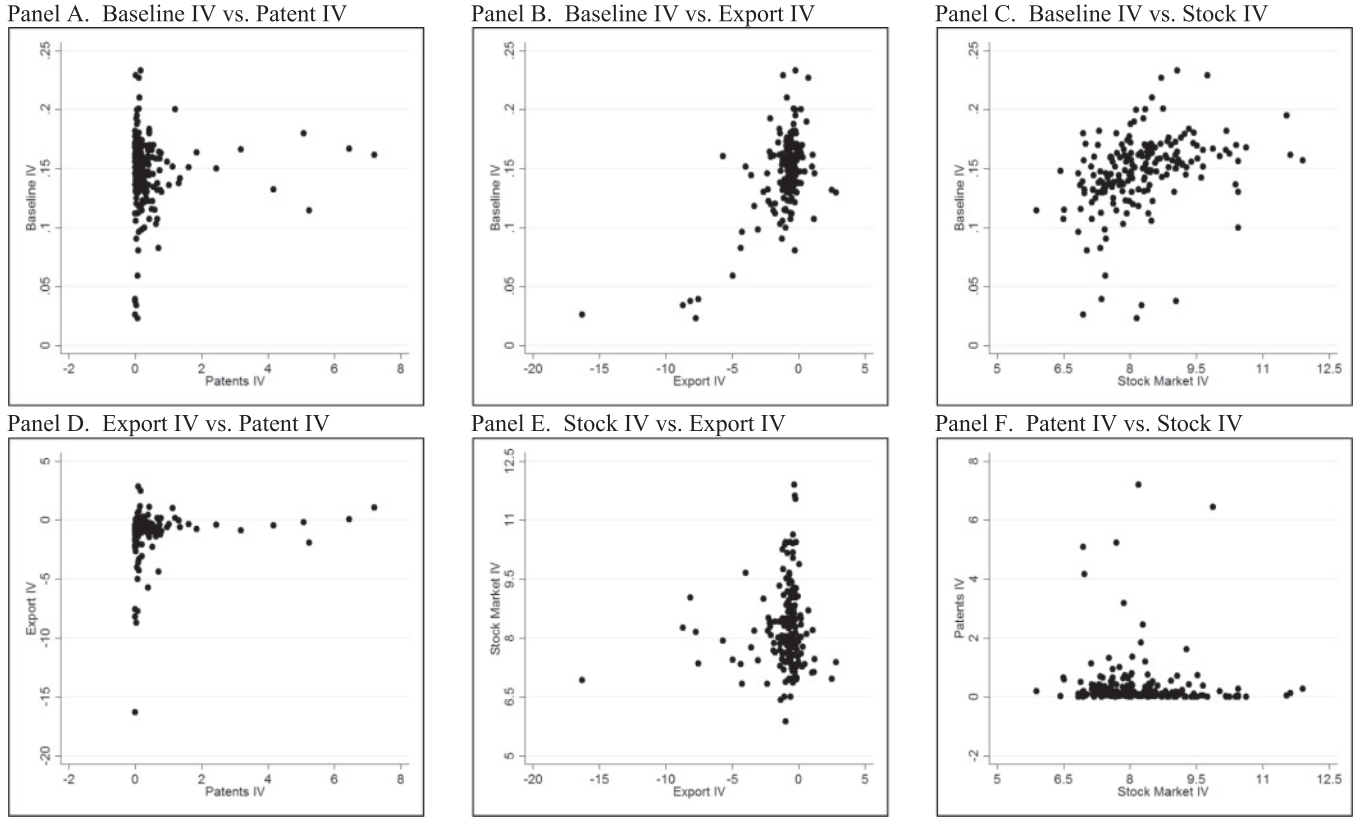
able. The estimated first-stage impact is reported at the bottom of the table, along with the F -statistic of the excluded instrument. Appendix table 2 reports corresponding OLS estimates.¹²

City employment responds substantially to local TFPR growth (panel A). A 1% increase in local TFPR is estimated to increase city employment by 2.38% in the medium run (column 1), by 4.16% in the long run (column 2), and by 4.03% in the longer run (column 3). These estimates suggest that it takes additional years for worker migration to respond to increased real wages and for housing construction to respond to increased demand, though this adjustment process was complete by 2000.

Panel B reports estimated impacts on annual earnings per worker. A 1% increase in local TFPR is associated with a

¹²The IV estimates are larger than the corresponding OLS estimates, which is consistent with the instrument reducing attenuation bias from measurement error in TFPR and downward bias from omitted variables. We also report cross-sectional OLS estimates that are generally larger in magnitude than the OLS estimates for changes in TFPR, which is also consistent with measurement error in TFPR.

FIGURE 2.—PAIRWISE CORRELATIONS BETWEEN ALTERNATIVE INSTRUMENTAL VARIABLES (BASELINE, PATENT, EXPORT, STOCK)



Each panel shows the pairwise correlation between two alternative instruments for predicting TFPR growth between 1980 and 1990: Baseline IV vs. Patent IV (coefficient 0.001, standard error 0.002, R -squared 0.002); baseline IV versus Export IV (coefficient 0.010, standard error 0.002, R -squared 0.358); baseline IV versus Stock Market IV (coefficient 0.011, standard error 0.002, R -squared 0.123); Export IV versus Patent IV (coefficient 0.221, standard error 0.079, R -squared 0.014); Stock Market IV versus Export IV (coefficient 0.041, standard error 0.030, R -squared 0.006); and Patent IV versus Stock Market IV (coefficient -0.076 , standard error 0.072, R -squared 0.006).


TABLE 2.—DIRECT EFFECT OF LOCAL TFPR GROWTH ON LOCAL EMPLOYMENT, EARNINGS, HOUSING COSTS (BASELINE IV)

	Medium-run Effect: Change from 1980 to 1990 (1)	Long-run Effect: Change from 1980 to 2000 (2)	Longer-run Effect: Change from 1980 to 2010 (3)
A. Log Employment	2.38*** (0.80)	4.16*** (1.26)	4.03*** (1.52)
B. Log Earnings	0.91*** (0.32)	1.45*** (0.47)	1.46*** (0.50)
C. Log Cost of Rent	0.98** (0.43)	1.47*** (0.46)	1.09** (0.48)
D. Log Home Value	1.74** (0.72)	2.46*** (0.78)	3.05*** (0.98)
E. Log Purchasing Power Renters	0.36** (0.18)	0.62** (0.26)	0.85*** (0.30)
Homeowners (Case A)	0.68*** (0.24)	1.11*** (0.37)	1.21*** (0.41)
Homeowners (Case B)	1.01*** (0.35)	1.60*** (0.51)	1.57*** (0.54)
First-Stage Coefficient	0.80*** (0.17)	0.80*** (0.17)	0.80*** (0.17)
Instrument F -statistic	23.64	23.64	23.64


Columns 1 to 3 report estimates from equations (1) to (3) in the text, respectively. Entries are the estimated coefficient on the change in city TFPR from 1980 to 1990. In column 1, the dependent variables are in changes from 1980 to 1990. In columns 2 and 3, the dependent variables are in changes from 1980 to 2000 (column 2) and in changes from 1980 to 2010 (column 3). In each column, we instrument for changes in city TFPR using the predicted change in TFPR, based on our baseline instrument. The corresponding first-stage estimate is reported in the row at the bottom of the table, with the associated F -statistic on the excluded instrument. In all specifications, the sample is our balanced sample of 193 MSAs. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

FIGURE 3.—SPATIAL DISTRIBUTION OF INSTRUMENTAL VARIABLES


Panel A. Baseline IV




Panel B. Patent IV



Panel C. Export IV



Panel D. Stock Market IV



For each indicated instrument, each panel shows the geographic variation in predicted TFPR growth from 1980 to 1990. Darker-shaded MSAs correspond to larger values of the instrument (and larger predicted growth in TFPR), with MSAs grouped into 10 equal-sized bins.

0.91% increase in earnings in the medium run (column 1), a 1.45% increase in earnings in the long run (column 2), and a 1.46% increase in earnings in the longer run (column 3). These estimated magnitudes are reduced-form effects of TFPR growth and, in particular, can be greater than 1 when worker in-migration and increased economic activity generate agglomeration spillovers (as in Greenstone, Hornbeck, & Moretti, 2010).

The estimated impacts on earnings are economically substantial. Given that real TFPR increased by 5.3% between 1980 and 1990 in the average city (appendix table 1), the IV estimates suggest that TFPR growth increased local earnings of full-time workers in the average sample city by 4.8% from 1980 to 1990, by 7.7% from 1980 to 2000, and by 7.7% from 1980 to 2010.

We expect increases in local housing costs to mitigate some portion of the estimated increases in local nominal earnings, given increases in employment that create additional demand for housing. Indeed, table 2 shows that increases in local TFPR are associated with substantially higher housing costs. A 1% increase in local TFPR leads to a 0.98% medium run increase in rental costs, a 1.47% long-run increase in rental costs, and a 1.09% longer-run in-

crease in rental costs (panel B). The corresponding effects on home values are somewhat larger, 1.74% to 3.05% (panel C), which suggests some expectation of future increases in rental costs. The estimated long-run effect on employment which are relative to rent is 2.8, and relative to home value is 1.7, similar magnitudes to the median housing elasticity (2.1) and mean housing elasticity (2.3) of sample MSAs with housing supply elasticity data from Saiz (2010).

We test whether the estimated impact on housing costs is larger in cities with a more inelastic housing supply, as a validation exercise suggested by previous empirical research (Glaeser & Gyourko, 2005; Gyourko, 2009; Saiz, 2010). We find that the estimated impact on local housing costs is indeed somewhat greater in cities with a more inelastic housing supply (appendix table 3). A 1% increase in local TFPR leads to a 2.3% long-run increase in rents in cities with below-average housing elasticity and to a 1.2% increase in rents in cities with above-average housing elasticity.

We note that growth in TFPR may reflect a variety of factors, including both technological improvement and increases in demand that raise prices. The impacts of local TFPR growth operate through increases in local labor demand, similar to increases in quantity TFP growth (TFPQ),

and so for our purposes, the important distinction between TFPR and TFPQ is not central and the empirical analysis draws on both sources of variation in TFPR.¹³

B. Direct Effects on Purchasing Power

The estimated increases in local earnings (table 2, panel B) are partly mitigated by increases in housing costs (panel C), but it remains unclear how much local TFPR growth affects local purchasing power (defined as increases in earnings net of increases in local cost of living). We use a measure of local cost of living that follows the BLS method for measuring the nationwide CPI but adapted to vary at the city level.

Renters. The effect of local TFPR growth on local purchasing power is conceptually straightforward for renters. In cities with TFPR growth, renters pay increased housing rents that reduce their purchasing power in proportion to the importance of housing as a share of total expenditures. Renters also pay increased costs of other nontradable goods, which reduces their purchasing power in proportion to the importance of nontradable goods as a share of total expenditures.

Thus, we define the effect on renter purchasing power as the percent increase in local earnings (panel B) minus the properly weighted percent increase in local rent (table 2, panel C) and the properly weighted percent increase in local prices of nonhousing, nontradable goods. The proper weights correspond to the share of total expenditures that is spent on housing and nonhousing, nontradable goods, respectively. We derive this expression in appendix C and address the important data limitation that there are not high-quality city-level data on the local price of nonhousing, nontradable goods for most cities in our time period. We follow the approach of Moretti (2013) to impute changes in prices of nontradables based on changes in rents.¹⁴ In practice, this means that we estimate the impact on renters' purchasing power as the estimated impact on log earnings minus 0.56 times log rent.

Panel E of table 2 reports that a 1% increase in local TFPR increases renters' purchasing power by 0.36% in the medium run (column 1) and by 0.62% in the long run (column

2). Purchasing power increases for renters because nominal earnings increase by more than the weighted increase in cost of living. Comparing the increase in purchasing power to the increase in nominal earnings, however, renters lose roughly two-thirds of their earnings increases to higher costs for housing and other local goods and services.

Homeowners. For workers who owned their home prior to the increase in TFPR, estimating changes in purchasing power is conceptually more complicated because it depends on how one accounts for the increase in home equity value. A large literature has examined the financial consequences of homeownership, but we are not aware of a widely accepted measure of the effect of housing prices changes on homeowners' purchasing power (Sinai & Souleles, 2005; Campbell & Cocco, 2007; Attanasio et al., 2009; Buiter, 2010; Mian, Rao, & Sufi, 2013; Ströbel & Vavra, 2019; Berger et al., 2018). Thus, we provide two bounds for how TFPR growth affects homeowners' purchasing power (see appendix C for more details). In one extreme case (case A), we consider a homeowner whose purchasing power is insulated only from increases in local housing rents.¹⁵ In case A, we measure the impact on homeowners' purchasing power as the estimated increase in earnings minus the properly weighted increase in the cost of nontradable goods (as calculated for renters, above).¹⁶ At the other extreme (case B), we consider a homeowner who is able to consume the wealth created by increased home value. In case B, we measure the impact on the homeowners' purchasing power as the estimated increase in earnings, plus the properly weighted increase in rental return on the home, minus the properly weighted increase in the cost of nontradable goods.¹⁷

Panel E reports that when homeowners are insulated from rising housing costs (case A), a 1% increase in local TFPR is associated with a 0.68% increase in purchasing power in the medium run (column 1) and a 1.11% increase in purchasing power in the long run (column 2). These increases in purchasing power are almost twice as large as the increases in purchasing power for renters, who face increased housing costs. The gains to homeowners are substantially larger when homeowners benefit from the increase in housing costs (case B). We conclude that local productivity growth

¹³Increases in physical productivity may also decrease goods prices and benefit consumers, whereas increases in goods prices (and TFPR) would induce welfare losses through the product market and knock-on effects that operate through input-output linkages with other industries. However, these effects would be largely national in scope for nationally traded products like manufacturing output and, thus, not affect our local analysis. Effects through input-output linkages are perhaps more local due to the more localized nature of supply chains, though this local effect would be reflected in the wages and employment effects that we estimate.

¹⁴Moretti (2013) infers how the prices of nonhousing, nontradable goods increase in a city along with increases in the cost of housing. He estimates that a 1% increase in the local rental price of housing is associated with a 0.35% increase in the local prices of other goods. Since the housing share of total expenditures is 0.33 (U.S. Bureau of Labor Statistics, 2000), we calculate the estimated impact on renters' purchasing power to be the estimated impact on log earnings minus 0.56 times log rent, where $0.56 = 0.33 + 0.35 \times (1 - 0.33)$. See appendix C for more details.

¹⁵This homeowner does not pay higher out-of-pocket housing costs when local housing prices increase, but the homeowner does pay a higher user cost for living in the home that is equal to the increased annual rental return on the home. The homeowner also faces increased local prices for nontradable goods, similar to renters.

¹⁶In practice, the change in purchasing power is measured as the estimated impact on log earnings minus 0.23 times log rent.

¹⁷In practice, the change in purchasing power is measured as the estimated impact on log earnings plus 0.10 times log rent (where $0.10 = 0.33 - 0.23$). This calculation assumes that homeowners can consume in perpetuity the annual return associated with increased housing rents in their city (i.e., the percent increase in housing rents multiplied by a 0.33 expenditure share on housing). We assume that homeowners can consume the increase in housing rents that would have been faced by renters of their home. Because homeowners' annual housing rents are unobserved, we assume homeowners and renters in the same city spend the same share of consumption on housing.

TABLE 3.—DIRECT EFFECT OF LOCAL TFPR GROWTH ON LOCAL EMPLOYMENT, EARNINGS, HOUSING COSTS (ADDITIONAL IVs)

	Medium-run Effect: Change from 1980 to 1990				Long-run Effect: Change from 1980 to 2000				Longer-run Effect: Change from 1980 to 2010			
	Stock IV (1)	Export IV (2)	Patent IV (3)	4 IVs Combined (4)	Stock IV (5)	Export IV (6)	Patent IV (7)	4 IVs Combined (8)	Stock IV (9)	Export IV (10)	Patent IV (11)	4 IVs Combined (12)
A. Log Employment	2.20*** (0.78)	3.94*** (1.44)	0.66 (0.82)	1.90*** (0.57)	2.92** (1.21)	5.89*** (2.25)	1.71 (1.46)	3.35*** (0.87)	3.47* (1.81)	5.31** (2.58)	1.31 (1.64)	3.32*** (1.10)
<i>p</i> -value of over-id test				0.32				0.47				0.66
B. Log Earnings	1.20*** (0.38)	1.53*** (0.56)	1.11** (0.53)	0.94*** (0.27)	1.72*** (0.56)	2.27*** (0.81)	2.08* (1.14)	1.54*** (0.43)	2.04*** (0.75)	1.85** (0.72)	2.22* (1.33)	1.67*** (0.51)
<i>p</i> -value of over-id test				0.25				0.31				0.45
C. Log Cost of Rent	1.75*** (0.56)	1.72*** (0.57)	2.13** (1.03)	1.26*** (0.40)	2.25*** (0.69)	2.13*** (0.68)	1.90* (1.05)	1.63*** (0.45)	2.24** (0.93)	1.47** (0.70)	2.54* (1.35)	1.52*** (0.53)
<i>p</i> -value of over-id test				0.10				0.16				0.20
D. Log Home Value	3.03*** (1.07)	2.98*** (0.90)	3.73** (1.87)	2.22*** (0.69)	2.45** (1.12)	3.55*** (0.95)	2.86* (1.49)	2.44*** (0.74)	4.41*** (1.65)	3.81*** (1.24)	4.33* (2.29)	3.45*** (1.01)
<i>p</i> -value of over-id test				0.14				0.48				0.49
E. Log Purchasing Power												
Renters	0.22 (0.20)	0.57* (0.31)	-0.09 (0.22)	0.24* (0.14)	0.46 (0.30)	1.08** (0.46)	1.02 (0.63)	0.63*** (0.24)	0.79* (0.41)	1.02** (0.40)	0.79 (0.77)	0.81*** (0.29)
<i>p</i> -value of over-id test				0.47				0.48				0.96
Homeowners (case A)	0.80*** (0.28)	1.13** (0.44)	0.62* (0.33)	0.65*** (0.19)	1.20*** (0.43)	1.78*** (0.66)	1.64* (0.92)	1.17*** (0.34)	1.53*** (0.58)	1.51*** (0.58)	1.63 (1.07)	1.32*** (0.41)
<i>p</i> -value of over-id test				0.43				0.38				0.67
Homeowners (case B)	1.38*** (0.43)	1.70*** (0.61)	1.32** (0.63)	1.07*** (0.30)	1.94*** (0.62)	2.49*** (0.88)	2.27* (1.24)	1.71*** (0.47)	2.27*** (0.83)	1.99** (0.78)	2.47* (1.45)	1.82*** (0.56)
<i>p</i> -value of over-id test				0.21				0.29				0.40
First-Stage Coefficient	0.021*** (0.006)	0.008** (0.003)	0.016** (0.007)		0.021*** (0.006)	0.008** (0.003)	0.016** (0.007)		0.021*** (0.006)	0.008** (0.003)	0.016** (0.007)	
Instrument <i>F</i> -statistic	13.34	6.31	5.86	9.47	13.34	6.31	5.86	9.47	13.34	6.31	5.86	9.47

The estimates correspond to those in table 2, using alternative instrumental variables. Columns 1, 5, and 9 use an instrument based on stock market returns. Columns 2, 6, and 10 use an instrument based on increased exposure to export markets. Columns 3, 7, and 11 use an instrument based on patenting activity. Columns 4, 8, and 12 use all four instrumental variables in combination, and below each estimate we report the *p*-value of the over-identification test (Hansen *J* statistic). Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

benefits local workers in large part through the housing market rather than through the labor market.

C. Additional Instrumental Variables (IVs)

Table 3 reports the estimated impacts of local TFPR growth using the stock market IV (columns 1, 5, 9), export IV (columns 2, 6, 10), and patent IV (columns 3, 7, 11) in the same form as the baseline IV estimates in table 2. The bottom row reports the estimated first-stage coefficients, along with the *F*-statistic on the excluded instruments. The estimated impacts fluctuate somewhat across specifications, and the alternative instruments have less power than our baseline instrument, but the pattern of results is generally consistent. Combining the use of all four instruments, columns 4, 8, and 12 report similar estimates as our baseline IV specifications. Overidentification tests fail to reject that the different instruments are yielding statistically different estimates. That is, despite drawing on identifying variation from different cities and industries experiencing different shocks, the instrumental variable estimates do not yield statistically different estimates of how local TFPR growth directly affects local economic outcomes. In particular, the stock market instrument isolates variation in TFPR growth that would be largely unanticipated, and these estimates suggest that our baseline estimates were not skewed by the anticipation of TFPR growth.

Given the similarity in the long-run estimates (1980 to 2000) and longer-run estimates (1980 to 2010) in tables 2 and 3, the remainder of the paper focuses on the long-run estimates and reports medium-run estimates as a point of comparison.

D. Direct Effects, by Worker Education

Table 4 reports our baseline IV estimates separately by worker education group.¹⁸ We estimate larger impacts of local TFPR growth on the employment of more educated workers, particularly in the long run (panel A). By contrast, we estimate larger impacts on earnings of less educated workers (panel B). These estimates are consistent with the notion that more educated workers are more geographically mobile in response to local economic shocks (Bound & Holzer, 2000; Wozniak, 2010; Malamud & Wozniak, 2012; Diamond, 2016; Notowidigdo, 2020). As a consequence, local more educated workers benefit less from local TFPR growth (in percentage terms).¹⁹

¹⁸Appendix table 4 reports the corresponding OLS estimates, though the OLS estimates are difficult to interpret.

¹⁹These differential effects on worker earnings are not undone by differential effects on housing costs by worker education group. There is some indication of lower impacts on rents and home values for more educated workers (panels C and D), but the impact on renters' and

TABLE 4.—DIRECT EFFECTS OF LOCAL TFPR GROWTH, BY EDUCATION LEVEL

	Medium-run Effect: Change from 1980 to 1990 (2SLS)				Long-run Effect: Change from 1980 to 2000 (2SLS)			
	College (1)	Some College (2)	High School or less (3)	Difference: (1) - (3)	College (5)	Some College (6)	High School or less (7)	Difference: (5) - (7)
A. Log Employment	2.79** (1.13)	2.60*** (0.73)	2.31*** (0.78)	0.48 (0.66)	5.82*** (1.88)	4.88*** (1.25)	3.23*** (1.15)	2.58*** (1.16)
B. Log Earnings	0.60** (0.24)	0.67*** (0.26)	1.12*** (0.30)	-0.52*** (0.20)	0.87** (0.34)	1.06*** (0.33)	1.23*** (0.35)	-0.36 (0.29)
C. Log Cost of Rent	0.55 (0.44)	1.02*** (0.38)	1.08** (0.47)	-0.53** (0.27)	1.01* (0.53)	1.50*** (0.40)	1.48*** (0.47)	-0.47 (0.39)
D. Log Home Value	1.59*** (0.58)	1.69** (0.74)	1.99*** (0.77)	-0.40* (0.30)	1.83*** (0.59)	2.14*** (0.71)	2.54*** (0.77)	-0.70*** (0.31)
E. Log Purchasing Power Renters	0.30 (0.22)	0.10 (0.13)	0.51*** (0.16)	-0.22 (0.25)	0.31 (0.20)	0.22 (0.15)	0.40** (0.18)	-0.09 (0.28)
Homeowners (Case A)	0.48** (0.20)	0.43** (0.18)	0.87*** (0.21)	-0.39** (0.21)	0.64** (0.25)	0.72*** (0.25)	0.89*** (0.26)	-0.25 (0.26)
Homeowners (Case B)	0.66** (0.27)	0.77*** (0.29)	1.23*** (0.34)	-0.57*** (0.20)	0.97** (0.38)	1.21*** (0.37)	1.38*** (0.39)	-0.40* (0.31)

Columns 1 to 3 report estimates that correspond to those in column 1 of table 2, but separately by education group: completed 4 years of college or more (column 1), completed between 1 and 3 years of college (column 2), and completed 12 years of education or fewer (column 3). Column 4 reports the difference between columns 1 and 3. Columns 5 to 8 report analogous estimates for the long-run effect by education group, corresponding to the estimates in column 2 of table 2. All entries are based on the baseline IV. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Panel E reports increases in the purchasing power of less educated renters in the medium run and long run, and somewhat smaller and statistically insignificant increases in the purchasing power of more educated renters. Local college-educated workers receive more substantial increases in purchasing power if they were homeowners prior to the TFPR shock, however, and are thereby insulated from increased housing costs or otherwise benefit from increased home values.

These results suggest that productivity growth reduces inequality at the local level, in both nominal terms and adjusted for local cost of living. College-educated workers appear very responsive to local TFPR shocks, whereas less educated workers are less responsive, and the greater employment responses for more educated workers appear to dampen the local economic gains for more educated workers. In the context of a Rosen-Roback model, where workers have idiosyncratic preferences for location, this would be the case if worker preferences for locations are relatively more important for less educated workers.²⁰

Figure 4 shows the differential responsiveness of workers by education group and its relationship with the local college earnings premium. Panel A shows a decreasing relationship between the change in local college earnings premium and predicted growth in local TFPR (using our baseline instrument). Panel B shows an increasing relationship between the change in employment share of college workers and predicted growth in local TFPR. This figure summarizes

homeowners' purchasing power is generally lower for college-educated workers (panel E).

²⁰This greater "preference" for locations among less educated workers could reflect a number of factors, including greater reliance on local family networks, greater benefits from local safety nets, or fixed costs of moving.

the intuition for how local TFPR growth decreases local inequality due to spatial mobility: following an increase in labor demand from local TFPR growth, a relative increase in the supply of more educated workers contributes to a decline in the education earnings premium.²¹

Table 5 reports estimated impacts of local TFPR growth on the distribution of local 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 lower portion of the distribution (panel C). Panel A, row 1 reports that increased local TFPR is associated with substantial declines in local earnings inequality. The estimated magnitude implies that a 1% increase in local TFPR reduces the 90-10 earnings gap by 0.632%, or that earnings at the 10th percentile increase by 0.632% more than earnings at the 90th percentile. This impact on inequality occurs at the upper portion of the distribution (panel B), whereas there is little impact on earnings inequality at the lower portion of the distribution (panel C).²² These effects are larger in the long run, with a 1% increase in local TFPR reducing the 90-10 earnings gap by 0.998% and reducing the 90-50 earnings gap by 0.930%.


One way to interpret the economic magnitude of our estimated effects is to relate the estimated impacts of TFPR to cities' elasticity of local labor supply. Local labor supply reflects how many workers are willing to live in a city for a given wage. Consider appendix figure 3, in which point 1

²¹Much of the literature on technological change and wage inequality has focused on the degree of skill bias in technological change, but we emphasize that even skill-neutral changes in local TFPR can differentially affect workers with different levels of education if they have different levels of geographic mobility.


²²Appendix table 5 reports the corresponding OLS estimates.

FIGURE 4.—LOCAL TFPR GROWTH AND LABOR MARKET OUTCOMES BY EDUCATION GROUP

Panel A. Change in College Earnings Premium



Panel B. Change in Share of College Workers



Panel A plots the change in city-level college earnings premium from 1980 to 1990 (log earnings of workers with four years of college education – log earnings of workers with no college education) against predicted local TFPR growth from 1980 to 1990 (based on our baseline instrument). The estimated coefficient is -0.495 (0.183). Panel B plots the change in city-level share of college workers with estimated coefficients of 0.108 (0.072). Circle sizes reflect MSA manufacturing output.

represents the equilibrium wage (w_1) and equilibrium employment (N_1) in a city before an increase in TFPR. An increase in TFPR then shifts local labor demand out from $D(TFP_1)$ to $D(TFP_2)$ and point 2 reflects the new equilibrium. By shifting labor demand, the TFPR shock identifies the slope of the function. The inverse elasticity of local labor supply is given by the ratio of the percent increase in earn-

ings over the percent increase in employment. When this ratio is smaller, the supply of labor to this city is more elastic and the supply curve is flatter. This reflects workers being more willing to move from other cities (without requiring much higher wages), as well as the housing stock being more able to adjust upward (without requiring much higher housing prices).

TABLE 5.—DIRECT EFFECTS OF LOCAL TFPR GROWTH ON LOCAL INEQUALITY

	Medium-run Effect: Change from 1980 to 1990 (1)	Long-run Effect: Change from 1980 to 2000 (2)
A. 90/10 Centile Difference in Log Earnings	-0.632*** (0.225)	-0.998** (0.420)
B. 90/50 Centile Difference in Log Earnings	-0.574*** (0.222)	-0.930*** (0.320)
C. 50/10 Centile Difference in Log Earnings	-0.058 (0.236)	-0.068 (0.292)

Column 1 reports estimates analogous to those reported in column 1 of table 2 (and column 2 reports estimates analogous to those reported in column 2 of 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). All entries are based on the baseline IV. Robust standard errors are reported in parentheses. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

From table 2, the estimated long-run impact on earnings (1.45), divided by the estimated long-run impact on employment (4.16), implies a long-run inverse elasticity of 0.35. This number reflects a relatively elastic local labor supply, indicating that in the long run, the U.S. labor force is fairly willing in this period to relocate to cities with better labor markets.²³

The local labor supply of college graduates is much more elastic than the local labor supply of high school graduates. The estimates by education group, from table 4, imply an inverse elasticity of 0.15 for college graduates and 0.38 for high school graduates.

We have been interpreting the estimated increases in employment as additional workers moving into the city, though increased employment could also reflect increased labor supply of existing city residents. Consistent with migration explaining most of the employment effect, we find that TFPR increases the level of employment and level of population by similar amounts (appendix figure 4).

E. Multiplier Effect on the Nonmanufacturing Sector

Increases in manufacturing TFPR have a direct effect on the manufacturing sector, but also indirectly affect the local nonmanufacturing sector. Wage and employment growth in manufacturing increase the demand for local nontraded goods and services, and therefore employment in nonmanufacturing sectors (Moretti, 2010). The extent to which nonmanufacturing sectors are affected is informative about how much policies directed at the manufacturing sector might influence the broader local economy. Indeed, policy efforts to support local manufacturing are often justified by policy-makers on these grounds.

Appendix table 6 reports that employment responds similarly in the manufacturing and nonmanufacturing sectors. These increases reflect a combination of in-migration and movement between sectors. We compute the implied multiplier effect of the manufacturing sector on the local non-

manufacturing sector, defined as the number of additional nonmanufacturing jobs created for each additional manufacturing job generated by TFPR gains.²⁴ From an increase in manufacturing TFPR that creates one manufacturing job, panel B reports an implied increase of 1.62 nonmanufacturing jobs. This estimate is consistent with estimates by Moretti (2010) based on a similar time horizon. 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, there is an implied increase of 2.21 nonmanufacturing jobs.

F. Alternative Specifications and Robustness

Pretrends in economic outcomes. We can extend the outcome data back to 1970 for 110 cities of our main sample of 193 cities and estimate the relationship between 1970-to-1980 outcome changes and instrumented changes in TFPR from 1980 to 1990. Appendix table 7 shows similar pretrends in employment and negative pretrends in wages and housing costs (column 1). The estimates are similar to our baseline estimates when controlling directly for 1970-to-1980 changes in the outcome variable (columns 2 and 3).

Serial correlation in TFPR. City TFPR growth from 1980 to 1990 is not strongly correlated with city TFPR growth from 1990 to 2000, with a slight negative relationship that is statistically insignificant. Appendix table 8 reports similar long-run estimates when controlling for changes in TFPR from 1990 to 2000 (column 1) and instrumenting for this later change in TFPR with an analogous instrument for that later period (column 2). City TFPR growth from 1990 to 2000 itself has generally smaller and statistically insignificant effects, but the first stage is notably less robust for

²³This elasticity is higher than that estimated by Beaudry, Green, and Sand (2014).

²⁴Local manufacturing TFPR growth may reflect local economic shocks that directly affect local nonmanufacturing sectors, but our IV estimates use variation in local manufacturing TFPR that is induced by national shocks to manufacturing industries that are less clearly related to local sources of TFPR growth in other sectors. An additional identification assumption here is that local nonmanufacturing sector growth is not otherwise associated with predicted changes in local manufacturing TFPR.

TFPR changes from 1990 to 2000.²⁵ The average impacts of TFPR growth may vary across time periods and contexts, based on where TFPR growth occurs and the characteristics of those places along with how TFPR growth translates into local labor demand. Indeed, we estimate that the effects of TFPR growth are concentrated in larger cities with less impact of TFPR growth on earnings and rents in smaller cities (appendix table 9, columns 1 and 2). The estimated effects of TFPR growth are more similar for cities that were previously growing at faster or slower rates (columns 3 and 4). Column 3 of appendix table 8 reports estimates from a long difference specification, regressing outcome changes from 1980 to 2000 on TFPR changes from 1980 to 2000 and instrumenting with the predicted change in TFPR from 1980 to 2000. The long difference specification may not reflect long-run effects, however, as changes in TFPR could occur any time between 1980 and 2000 and the estimated magnitudes are more similar to the medium-run estimates in table 2.

Spatial correlation in TFPR. Appendix table 8 reports similar estimates when controlling for TFPR growth in cities within 500 miles, 250 miles, and 100 miles (columns 4–6), instrumenting using predicted TFPR growth in those cities based on their industry shares and industry-level TFPR growth. These specifications also effectively control for regional industry concentration, exploiting variation in relative local industry concentration within that particular city. Cross-city correlations may also affect the statistical inference. To allow for potential geographic correlation among nearby MSAs, appendix table 10 reports our baseline estimates when clustering the standard errors by state (42 clusters, columns 2 and 4) or contiguous MSA groupings (114 clusters, columns 1 and 3). Appendix table 11 reports estimates using the procedure developed by Adao, Kolesar, and Morales (2019) to allow for correlation across MSAs with similar baseline industry shares that vary by instrument in predicting city TFPR growth. The inference remains similar to our baseline estimates, consistent with the substantial spatial variation across cities with different industry shares.

Additional controls. The baseline estimates are not sensitive to controlling for cities' total manufacturing share (appendix table 12, columns 1 and 5), as suggested by Borusyak, Hull, and Jaravel (2021) for similar research designs.²⁶ Columns 2 and 6 report estimates controlling for cities' employment share in 1980 in broad industry categories outside of manufacturing. Columns 3 and 7 report es-

timates when controlling for cities' 1980 employment share in the oil and gas industry, which experienced particularly negative shocks in the 1980s (table 1).

Given the estimated increases in local employment following local TFPR growth, one question is whether changes in worker composition are driving the estimated increases in annual earnings. Columns 4 and 8 of panel B in appendix table 12 report estimated impacts using individual-level data to condition on worker characteristics: age, age-squared, education, race, and gender. Panels C and D report similar estimated impacts on housing costs when using individual-level data to condition on physical characteristics of the home: the number of rooms and number of bedrooms, whether the home is part of a multi-unit structure, and the presence of a kitchen or plumbing.²⁷ These specifications are not our preferred models, however, because the changes in worker composition are endogenous and conditioning on endogenous responses to local TFPR growth can introduce bias.

Contamination of control group. When local employment increases in cities that experience relatively greater TFPR growth, some of those workers are drawn from other sample cities that make up the control group. The estimated relative effects on employment would then be biased upward, as the average comparison city is negatively affected, indirectly increasing wages and decreasing housing costs in comparison cities. We expect this contamination bias to be small, however, because there is little average indirect effect on comparison cities.²⁸ Some sample MSAs are more closely linked with particular other sample MSAs, whereby one MSA receives more than 10% or 5% of its migrants from that other MSA. Appendix table 13 reports estimates when aggregating the data from these MSAs and treating them as one observation (columns 1, 2, 5, 6). Columns 3 and 7 report medium-run and long-run estimates when combining contiguous MSAs into one MSA. Columns 4 and 8 report estimates when omitting region fixed effects, such that the comparison cities are all other sample MSAs.

Because each city is a small share of the total labor market, the indirect effects are spread across many cities and there is a negligible indirect impact on the average control city. In considering the sum of these small indirect effects on each other city, however, the total indirect effect may be substantial.

V. Indirect Effects of Local TFPR Growth

Estimates from section IV report how a local TFPR shock affects employment, wages, and housing costs in the city where the shock occurs, relative to other cities. These

²⁵In particular, there is no first stage for TFPR growth from 1990 to 2000 for cities with less elastic housing supply, which we have viewed as a way of validating the estimates, and the 1990–2000 changes in TFPR are more driven by outlier cities.

²⁶Our baseline estimates also use the share of manufacturing activity in each manufacturing industry, rather than the share of total activity in each manufacturing industry, to focus on how industry-wide manufacturing TFPR shocks may differentially affect manufacturing TFPR growth across cities.

²⁷Panel E then reports impacts on purchasing power including both sets of control variables.

²⁸Our main estimates include region fixed effects, but there is substantial cross-region migration and nonsample MSAs such that the average sample MSA receives 35% of its migrants from other sample MSAs in the same region (based on 1975–1980 migrant flows in U.S. Census data) and those migrants would be dispersed among sample MSAs within the region.

estimated direct effects on local outcomes are only part of the overall impact from a local TFPR shock, however, as the local TFPR shock also has indirect effects outside that particular city.

We propose a methodology for quantifying indirect effects generated through worker mobility. Our approach builds on the estimated direct effects, along with particular assumptions about the elasticity of labor demand and patterns of worker mobility.

Intuitively, local TFPR growth in one city (e.g., Houston) generates additional indirect impacts on labor markets and housing markets in other cities due to worker migration responses. Some migrants to Houston come from Dallas, which raises wages and lowers housing costs in Dallas given downward-sloping labor demand and an upward-sloping housing supply in Dallas. Dallas also experiences its own TFPR shock, as do other cities, but this indirect effect represents the pressure on labor markets and housing markets in Dallas from TFPR changes in Houston.

The magnitude of indirect effects depends on the magnitude of worker reallocation, and our estimates from section IV found substantial direct employment effects in response to local TFPR growth (particularly in the long run). We therefore expect the indirect effects to be substantial, and particularly large for more educated workers who are more mobile. Thus, we anticipate that local TFPR shocks will have different impacts on inequality at the aggregate level as compared to the direct effects on inequality at the local level.

For each sample city, we use our estimated direct effect on employment and data on city-to-city migration to estimate how TFPR growth in that city alone would induce employment changes in the other sample cities. We then use data on cities' elasticity of housing supply, along with an assumption about the elasticity of labor demand, to quantify the indirect effects on housing costs and worker earnings in these other sample cities. We sum these indirect effects from TFPR growth in each sample city and compare this magnitude to the estimated direct effects on sample cities. Specifically, we proceed in three steps.

Step 1. For each of the 193 sample cities c , we use estimates from section IV to calculate the number of workers drawn to city c from 1980 to 2000 based on its growth in TFPR from 1980 to 1990. This number is the product of city c 's growth in TFPR from 1980 to 1990, times the estimated long-run impact on employment (table 2, panel A, column 2), times city c 's baseline employment in 1980 (appendix table 1).

Step 2. Given an increase in workers in city c , we calculate the associated number of workers who would leave each of the other 192 cities o due to TFPR growth in city c . Because we do not observe where these workers would move from, in response to increasing TFPR in city c only, we use data on observed city-to-city migration rates to characterize typical cross-city migration links. As a baseline assumption, we assume that workers are drawn to city c from city o in

proportion to observed migration flows from 1975 to 1980 in the 1980 Census of Population.²⁹ For example, if Houston would have added 1,000 new workers between 1980 and 2000 (based on its TFPR gains from 1980 to 1990 and the estimated impact of local TFPR growth on local employment), and 5% of migrants to Houston were from Dallas from 1975 to 1980, then we calculate an induced decline of 50 workers in Dallas from TFPR growth in Houston (all else equal). As an alternative method for assigning migrant origins, we assume that workers moving to city c are drawn from all other locations in proportion to their size (which holds fixed the relative sizes of other cities). As another method, we assign the share of migrants to city c from city o using predicted migrant flows from an estimated gravity equation.³⁰

Step 3. Given the induced change in employment in each other origin city o , from TFPR growth in city c , we calculate the resulting pressure on housing costs and earnings in city o . For housing costs, we calculate the decline in households in city o , based on the decline in workers and the average number of workers per household in city o , and use the estimated city-level elasticities of housing supply from Saiz (2010).³¹ For earnings, our baseline calculations assume a constant elasticity of labor demand (-0.15). We also report estimates allowing for heterogeneity across cities in the elasticity of labor demand due to variation in city industry mix.³² We assume no agglomeration economies, whereby changes in city employment would affect city TFPR (and then affect city employment, and so on).

These three steps provide an estimate of how a local TFPR change in each city c indirectly affects wages and housing costs in each other city o through worker mobility. We sum these indirect effects across all cities o . We then sum the indirect effects from each city c and compare these to the direct effects on all cities c .

Appendix D illustrates this approach with the examples of local TFPR growth in Houston, San Jose, and Cincinnati.

²⁹We assume a closed economy without international migration, in which a fixed number of workers move across sample MSAs.

³⁰Drawing on a literature estimating gravity equations in migration flows, we regress city-to-city migrant flows between 1975 and 1980 on the log size of origin city o , log size of destination city size c , the log geographic distance between city o and city c , and the log economic distance between city o and city c (defined as the vectorial distance in the cities' industry employment shares).

³¹The estimated elasticities of housing supply from Saiz (2010) reflect the responsiveness of local house prices to local demand shocks, whereas our estimated impacts on purchasing power use the responsiveness of rental costs to local demand shocks. We estimate that rental costs are less responsive than house prices, as is typical in the literature, and so we scale the estimates from Saiz (2010) by the ratio of our estimated impacts on rental costs and housing prices (table 2, column 2, panels C and D) to obtain an elasticity of rental costs with respect to local demand. The resulting average elasticity is 2.7, weighting by worker population, such that a 1% decrease in workers would decrease rental costs by 0.37%.

³²We use data on labor shares by two-digit SIC industry and calculate industry-specific labor demand functions assuming the elasticity of labor demand is equal to 1 minus the labor share minus a flexible capital share (0.20). We then calculate city-level labor demand elasticities by weighting each industry based on its initial employment share. By comparison, our baseline calculation assumes a constant labor share of 0.65 and a flexible capital share of 0.20.

TABLE 6.—LONG-RUN DIRECT EFFECTS, INDIRECT EFFECTS, AND COMBINED EFFECTS OF LOCAL TFPR GROWTH

	Long-run Direct Effects on:				Long-run Indirect Effects on:				Total Effect (9)	Total % Effect (10)	Annual Total % Effect (11)	Robustness: Annual	
	Earnings (1)	Housing (2)	Nontradables (3)	Purchasing Power (4)	Earnings (5)	Housing (6)	Nontradables (7)	Purchasing Power (8)				Total % Effect (12)	Total % Effect (13)
A. Renters	3,823 (1,368)	1,286 (449)	900 (314)	1,636 (605)	919 (392)	-1,044 (447)	-731 (313)	2,693 (1,152)	4,329 (1,312)	11.2% (3.4%)	0.56% (0.17%)	0.58% (0.19%)	0.47% (0.14%)
B. Homeowners													
Case A	5,008 (1,807)	-	1,180 (415)	3,828 (1,392)	1,331 (566)	-	-948 (403)	2,279 (969)	6,107 (1,708)	11.3% (3.2%)	0.57% (0.16%)	0.57% (0.16%)	0.54% (0.16%)
Case B	5,008 (1,807)	1,685 (593)	1,180 (415)	5,514 (1,985)	1,331 (566)	-1,354 (576)	-948 (403)	924 (393)	6,438 (2,027)	12.0% (3.8%)	0.60% (0.19%)	0.59% (0.19%)	0.60% (0.19%)

Entries are the average per worker direct effects, indirect effects, and combined total effects of 1980 to 1990 TFPR growth on 1980 to 2000 changes in outcomes in 2017 dollars. Columns 1 to 3 report direct effects of TFPR growth on earnings, housing costs, and the cost of nonhousing, nontradable goods. Column 4 reports the direct effect on purchasing power. The effect on purchasing power for renters (panel A) is defined as column 1 - column 2 - column 3. For homeowners (panel B), the effect on purchasing power in case A is defined as column 1 - column 3; in case B, it is defined as column 1 + column 2 - column 3. Columns 5 to 7 report indirect effects of TFPR growth on earnings, housing costs, and the cost of nonhousing nontradable goods. Column 8 reports the indirect effect on purchasing power. Column 9 reports the total effect, defined as the sum of the direct effect and indirect effect. Column 10 expresses the total effect as a percent increase relative to 1980 average earnings (in 2017 dollars). Column 11 expresses these numbers in annual terms, dividing column 10 by 20. Columns 12 and 13 report robustness to alternative assumptions on mobility: in column 12, that migration flows from other sample cities in proportion to their population sizes; in column 13, that migration flows are based on predicted migration flows only (taking the predicted values from regressing 1975-1980 migrant flows on log origin city size, log destination city size, log geographic distance, and log economic distance). Robust standard errors are reported in parentheses.

For Houston, its TFPR growth alone would increase employment by 86,031 workers, earnings by \$1,490 per worker and rent by \$501 per worker (appendix table 14, panel A, column 1). This increased employment in Houston would draw some workers from Dallas (4,551), San Antonio (2,617), and Boston (374) among other places, with an average decline of 291 workers in other sample cities that is associated with a \$9 increase in earnings and a \$8 decline in rent, on average, for workers in other sample cities (appendix table 14, panel A, column 2). These indirect effects in each of the other cities are small, on average, but these indirect effects will be economically substantial when summed across all cities.

Our approach to calculating these indirect effects requires fewer assumptions than studies that identify general equilibrium effects of local shocks using the structure of a spatial equilibrium model.³³ Our analysis is more limited in scope, however, and considers only indirect effects stemming from worker mobility. There may exist other types of indirect effects in general equilibrium, such as on the price of traded goods or the returns to capital. Quantifying these other general equilibrium effects is outside the scope of this paper and requires stronger model assumptions.

VI. Combined Impacts of Local TFPR Growth

A. Direct Effects, Indirect Effects, and Combined Effects

Table 6 reports the long-run impact of local TFPR growth through direct effects (columns 1-4), indirect effects

(columns 5-8), and the combined effect on worker purchasing power (columns 9-13).³⁴

Panel A reports that local TFPR growth had substantial long-run direct effects on the average renter's earnings (\$3,823), housing costs (\$1,286), and costs of other local goods (\$900) in the cities directly hit by TFPR shocks.³⁵ The direct effect on purchasing power for renters (\$1,636) reflects increased cost-of-living offsetting two-thirds of the increase in earnings. Summing the indirect effects of local TFPR growth in each city, however, the average renter received a substantial further increase in earnings (\$919), a decrease in housing costs (-\$1,044), and a decrease in cost of other local goods (-\$731). These indirect effects contributed a net increase of \$2,693 in renters' purchasing power (column 8).

Indirect effects make up almost two-thirds of the combined \$4,329 increase in purchasing power for renters (column 9). For renters, most of the increase in local housing costs from increased local TFPR is offset by decreases in local housing costs from increased TFPR in other cities.³⁶ Along with indirect increases in earnings, the combined increase in renters' purchasing power reflects an 11.2% increase on 1980 earnings (column 10) or 0.56% annual increase from 1980 to 2000 (column 12). These numbers are

³³For recent examples on the general equilibrium effects of local productivity, see Caliendo et al. (2018) and Hsieh and Moretti (2019). See also recent complementary work on the general equilibrium effects of trade shocks (Caliendo, Dvorkin, & Parro, 2019; Adao, Arkolakis, & Esposito, 2023) and on mobility across areas (Monte, Rossi-Hansberg, & Redding, 2018).

³⁴We calculate the combined effects by summing the direct effects and indirect effects from local TFPR growth in each city, sum these effects across each city, and then divide by the total number of workers in sample cities. The standard error of the combined effect follows from the variance-covariance structure of the estimated direct effects and the estimated correlation across MSAs between the direct effects and indirect effects.

³⁵Following our discussion of impacts on purchasing power, we assume that the dollar cost of other local goods increases by 0.70 times the dollar increase in housing costs, which reflects a 0.35% increase in the cost of other goods from a 1% increase in housing costs, along with an expenditure share on other goods that is twice the expenditure share on housing (0.33).

³⁶These effects need not cancel, as the elasticity of housing supply varies across cities and so it matters which cities are experiencing local TFPR growth.

TABLE 7.—LONG-RUN DIRECT EFFECTS, INDIRECT EFFECTS, AND COMBINED EFFECTS OF LOCAL TFPR GROWTH BY WORKER EDUCATION GROUP

	Long-run Direct Effects on:				Long-run Indirect Effects on:				Total Effect (9)	Total % Effect (10)	Annual Total % Effect (11)	Robustness: Annual	
	Earnings (1)	Housing (2)	Nontradables (3)	Purchasing Power (4)	Earnings (5)	Housing (6)	Nontradables (7)	Purchasing Power (8)				Total % Effect (12)	Total % Effect (13)
A. Workers with College Education													
Renters	3,173 (1,311)	1,223 (692)	856 (485)	1,094 (136)	2,005 (503)	-1,430 (399)	-1,001 (279)	4,437 (816)	5,531 (824)	10.8% (1.6%)	0.54% (0.08%)	0.62% (0.16%)	0.54% (0.15%)
Homeowners													
Case A	4,514 (1,877)	-	1,219 (695)	3,294 (1,182)	3,505 (1,089)	-	-1,490 (412)	4,995 (1,139)	8,289 (1,618)	10.3% (2.0%)	0.52% (0.10%)	0.55% (0.16%)	0.54% (0.16%)
Case B	4,514 (1,877)	1,742 (993)	1,219 (695)	5,036 (2,175)	3,505 (1,089)	-2,128 (589)	-1,490 (412)	2,866 (1,294)	7,902 (2,499)	9.8% (3.1%)	0.49% (0.16%)	0.52% (0.20%)	0.52% (0.20%)
B. Workers with High School Education or Less													
Renters	2,853 (884)	1,156 (408)	809 (286)	889 (192)	521 (243)	-936 (257)	-655 (180)	2,111 (633)	3,000 (664)	8.5% (1.9%)	0.43% (0.09%)	0.45% (0.11%)	0.36% (0.08%)
Homeowners													
Case A	3,558 (1,108)	-	1,010 (359)	2,548 (749)	688 (320)	-	-796 (217)	1,484 (499)	4,031 (910)	8.8% (2.0%)	0.44% (0.10%)	0.44% (0.10%)	0.42% (0.10%)
Case B	3,558 (1,108)	1,443 (513)	1,010 (359)	3,991 (1,262)	688 (320)	-1,137 (310)	-796 (217)	347 (151)	4,337 (1,280)	9.4% (2.8%)	0.47% (0.14%)	0.47% (0.14%)	0.48% (0.14%)
C. Average Impacts by Worker Education													
Workers with College Education				3,204 (1,195)				4,089 (1,090)	7,293 (1,669)	10.3% (2.3%)	0.52% (0.12%)	0.56% (0.17%)	0.53% (0.17%)
Workers with High School Education				2,446 (724)				1,329 (410)	3,774 (902)	8.9% (2.1%)	0.44% (0.11%)	0.45% (0.11%)	0.42% (0.10%)

Panels A and B report estimates similar to table 6, but separately by worker education group. Panel C reports average impacts for each worker education group, weighting by the fraction of renters or homeowners (for homeowners, we take the average of case A and case B). Robust standard errors are reported in parentheses.

similar under alternative assumptions for worker migration flows (columns 12 and 13), which result in less concentrated migrant flows between particular cities compared to the observed migrant flows from 1975 to 1980 but yield similar estimates of total indirect effects.

Panel B reports impacts on homeowners. Compared to renters, homeowners receive larger direct effects on purchasing power because homeowners do not pay higher housing rents (case A) or even benefit from local increases in housing rents (case B).³⁷ Homeowners benefit less than renters from the indirect effects of TFPR growth, however, because of decreasing housing rents due to TFPR growth in other cities (columns 5–8). Gains for homeowners in some cities come at the expense of homeowners in other cities. For homeowners, only 26% of their combined increase in purchasing power comes from indirect effects (taking the average of case A and case B).

Renters and homeowners receive notably similar percent increases in purchasing power from local TFPR growth when including both direct effects and indirect effects

³⁷For this case B, as above, we assume that homeowners can consume the increase in housing rents that would have been faced by renters of their home. Homeowners' annual housing rents are unobserved, so we assume homeowners and renters in the same city spend the same share of earnings on annual housing rents. Homeowners also receive a larger increase in earnings than renters (column 1), largely because their baseline average earnings are higher. We assume that local TFPR growth has the same percent effect on local earnings of renters and homeowners, but the geographic distribution of homeowners and renters also matters because they may be disproportionately in cities that experience different changes in TFPR.

(columns 10 and 11). The estimated direct effects imply much larger purchasing power gains for homeowners, compared to renters, but this disproportionate benefit is entirely counterbalanced by the estimated indirect effects. For the average worker, taking a weighted average over renters and homeowners, 38% of the overall increase in workers' purchasing power occurs outside cities directly affected by local TFPR growth.³⁸ While TFPR growth in one city has small indirect effects on each other city, on average, the sum of these indirect effects is substantial and reshapes who benefits to what degree from local TFPR growth.³⁹

B. Combined Effects by Education Group

Table 7 reports the direct effects and indirect effects separately for more educated workers (panel A) and less educated workers (panel B). Panel C reports the average impact

³⁸For calculating this weighted average, the weights reflect the share of workers that are renters (33.6%) and homeowners (66.4%). For homeowners, we take the average of case A and case B.

³⁹TFPR growth in sample cities also generates indirect effects outside sample MSAs, raising wages and decreasing housing costs when workers are drawn to sample MSAs, though we do not have the data to quantify these effects along with the direct effects of TFPR changes in nonsample areas and their indirect effects on sample MSAs and nonsample areas. If smaller cities or rural areas experience less direct effect from local TFPR growth, as suggested by appendix table 9, then these areas would experience predominantly indirect effects when TFPR increases in sample MSAs that decrease housing costs and increase wages (particularly for more educated workers). The aggregate impacts would then skew more toward increases in worker real earnings.

by education group, averaging over renters and homeowners based on homeownership rates by education group.⁴⁰

There are similar annual percent increases in purchasing power for more educated workers (0.52%) as for less educated workers (0.44%) when summing the direct effects and indirect effects (panel C, column 11). The direct effects on purchasing power are only moderately higher, in levels, for more educated workers (table 7, panel C, column 4) because of the larger estimated percent gains for less educated workers (table 4).⁴¹ The indirect effects on purchasing power, however, are substantially higher for more educated workers (table 7, panel C, column 8).⁴² Because of higher geographic mobility among more educated workers, there are substantially greater indirect increases in earnings of more educated workers in both levels and percentage terms. Indirect effects make up 56% of the overall effect for more educated workers, compared to 35% of the overall effect for less educated workers (panel C, columns 8 and 9). While we estimated that local TFPR growth compresses local inequality, the presence of indirect effects causes local TFPR growth to have little effect on inequality by worker education (in percent terms).

TFPR shocks do have substantial redistributive effects across workers in different locations, however, by education group and homeownership status. Local TFPR shocks benefit local less educated workers more than local more educated workers, and they benefit more educated workers in other cities more than less educated workers in other cities. More educated workers benefit wherever local TFPR increases due to their greater geographic mobility, whereas less educated workers are more sensitive to TFPR shocks within their city. Local TFPR shocks also benefit local homeowners more than local renters, whereas these shocks benefit renters in other cities more than homeowners in other cities. These effects have important implications for the geographic distribution of gains from productivity growth, as well as who benefits from productivity growth within those areas.

C. Combined Effects by Location

The impacts of TFPR shocks are very different across the country. This is because TFPR growth is heterogeneous across locations, so the direct effects vary across cities, but

⁴⁰For homeowners, we take the average impact on purchasing power for case A and case B. We then calculate the weighted average impact within each education group, weighting by the fraction of workers who are renters or homeowners among college-educated workers (31.3% renters) and high-school-educated workers (34.6% renters).

⁴¹This is also despite a slightly higher share of homeowners among more educated workers, which increases the direct effect on purchasing power from local TFPR growth. These estimates also reflect the geography of TFPR shocks, which matters due to variation across cities in their share of more educated workers.

⁴²Note that we assume no imperfect substitution between more educated and less educated workers, as well as no externalities across workers. That is, when calculating indirect effects by education group, we assume that out-migration of more educated workers affects only more educated workers' earnings and that out-migration of less educated workers affects only less educated workers' earnings.

also because the indirect effects vary substantially across cities when cities are connected differentially to cities that experience different TFPR shocks. This means local TFPR shocks have important redistributive effects across space. For example, local TFPR growth in Houston benefits workers and landowners in Houston, benefits workers and renters in Dallas, and hurts landowners in Dallas.

These effects do not necessarily balance out over geographic space, as some cities are positioned to receive larger indirect effects independent of the magnitude of their own direct effects. Appendix figure 8 maps the substantial variation across cities in the direct effects, indirect effects, and combined effects for renters. Appendix figure 9 shows there is little inherent correlation between cities that receive large direct effects and cities that receive large indirect effects. Thus, while indirect effects magnify the direct effects of local TFPR growth, the indirect effects of TFPR growth elsewhere do not inherently compensate workers for the relative absence of direct effects in their city.⁴³ Even at the regional level, there remains substantial variation in the relative contribution of indirect effects, and workers' location matters substantially for the benefits they receive from productivity growth.

VII. Conclusion

We make two contributions. On a substantive level, we estimate who benefits when cities experience productivity growth. We find that the average U.S. worker benefited substantially from manufacturing TFPR growth, though these gains depend substantially on where workers live. A high-level view of average changes would mask substantial variation in benefits across areas and people. On a methodological level, we propose a new approach to estimate general equilibrium effects of local shocks.

We find that when a city experiences TFPR growth in manufacturing, local earnings increase, but in-migration also raises local housing costs. For workers who rent their home, increased earnings are in large part offset by the increased cost of living, while the benefits for homeowners are more substantial. Thus, at the local level, TFPR growth benefits the average local worker but many of the benefits come through the housing market rather than the labor market.

Local TFPR growth reduces local inequality. Local TFPR shocks have more impact on the earnings of local less educated workers than the earnings of local more educated workers. There is greater in-migration of more educated workers, consistent with more educated workers being more geographically mobile on average.

⁴³Appendix tables 15 and 16 divide cities based on the terciles of direct effects and indirect effects and list example cities that received large direct effects and large indirect effects (panel A), large direct effects and small indirect effects (panel B), small direct effects and large indirect effects (panel C), and small direct effects and small indirect effects (panel D). Example cities in the top group for renters are Binghamton, Charleston, New Orleans, and San Jose. Example cities in the bottom group are Dallas, St. Louis, Tulsa, and Youngstown.

Local TFPR growth also has important indirect effects on other cities, however, and these indirect effects are large enough to alter the ultimate incidence of local TFPR growth. We estimate that 38% of the overall increase in purchasing power for the average worker occurs outside cities directly affected by local TFPR growth. Neglecting these indirect effects, generated by worker mobility, would substantially understate the gains from local productivity growth and misstate the distributional consequences.

The indirect effects on worker earnings are substantially greater for more educated workers, who migrate more to cities with increasing TFPR, which increases inequality in other cities. The net percent impact on purchasing power is then similar across less educated and more educated workers, with less educated workers benefiting more locally and more educated workers benefiting more elsewhere.

The net impact on purchasing power is also similar for renters and homeowners, with homeowners benefiting more locally and renters benefiting more elsewhere. Due to these indirect effects, the impacts on landowners are largely a transfer from one location to another. The overall incidence of TFPR growth then falls mainly on workers, though workers' location matters substantially and especially so for less educated workers who are less geographically mobile.

From a methodological point of view, our approach to including indirect effects may be helpful to those seeking to estimate the effects of economic shocks by comparing areas that experience large shocks to areas that do not. Our findings indicate that when local shocks generate large migration responses, a substantial portion of the overall effects may be missed when focusing only on the direct effects. Including these indirect effects, even those indirect effects from worker mobility only, can yield qualitatively and quantitatively different conclusions. Our approach can be used in other contexts to gauge the magnitude of indirect effects in a reduced-form manner.

REFERENCES

- Adao, Rodrigo, Costas Arkolakis, and Federico Esposito, "General Equilibrium Effects in Space: Theory and Measurement," working paper (2023). <https://sites.google.com/site/tradao/research>
- Adao, Rodrigo, Michal Kolesar, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics* 134:4 (2019). 10.1093/qje/qjz025
- Aghion, P., A. Bergeaud, M. Lequien, and M. Melitz, "The Heterogeneous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports" (2020).
- Attanasio, Orazio P., Laura Blow, Robert Hamilton, and Andrew Leicester, "Booms and Busts: Consumption, House Prices and Expectations," *Economica* 76:301 (2009), 20–50.
- Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review* 103:6 (2013), 2121–2168. 10.1257/aer.103.6.2121
- Autor, David, David Dorn, Gordon H. Hanson, Gary Pisano, and Pian Shu, "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," *American Economic Review: Insights* 2:3 (2020). 10.1257/aeri.20180481
- Beaudry, Paul, David Green, and Ben Sand, "Spatial Equilibrium with Unemployment and Wage Bargaining: Theory and Estimation," *Journal of Urban Economics* 38:3 (2014), 34–70.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *Review of Economic Studies* 85:3 (2018), 181–213. 10.1093/restud/rdx060
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies* 89:1 (2021).
- Bound, John, and Harry J. Holzer, "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s," *Journal of Labor Economics* 18:1 (2000), 20–54. 10.1086/209949
- Buiter, Willem H., "Housing Wealth Isn't Wealth," *Economics: The Open-Access, Open-Assessment E-Journal* 4 (2010), 2010–2022. 10.5018/economics-ejournal.ja.2010-22
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro, "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica* 87 (2019), 741–8310.3982/ECTA13758
- Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte, "The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy," *Review of Economic Studies* 85:4 (2018), 2042–2096. 10.1093/restud/rdx082
- Campbell, John Y., and Joao F. Cocco, "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics* 54:3 (2007), 591–621. 10.1016/j.jmoneco.2005.10.016
- Card, David, "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics* 19:1 (2001), 22–64. 10.1086/209979
- Charles, Kerwin, Eric Hurst, and Matthew Notowidigdo, "The Masking of the Decline in Manufacturing Employment by the Housing Bubble," *Journal of Economic Perspectives* 40:1 (2018), 7–72.
- Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000," *American Economic Review* 106:3 (2016), 479–524. 10.1257/aer.20131706
- Diamond, Rebecca, and Enrico Moretti, "Where Is Standard of Living the Highest? Local Prices and the Geography of Consumption," working paper (2022).
- Glaeser, Edward L., and Joseph Gyourko, "Urban Decline and Durable Housing," *Journal of Political Economy* 113:2 (2005). 10.1086/427465
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review* 110:8 (2020). 10.1257/aer.20181047
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti, "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy* 118:3 (2010), 536–598. 10.1086/653714
- Gyourko, Joseph, "Housing Supply," *Annual Review of Economics* 1:1 (2009), 295–318. 10.1146/annurev.economics.050708.142907
- Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools," NBER working paper 8498 (2001).
- Hsieh, Chang-Tai, and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics* 11:2 (2019), 1–39. 10.1257/mac.20170388
- Jones, Charles I., and Peter J. Klenow, "Beyond GDP? Welfare across Countries and Time," *American Economic Review* 106:9 (2016), 2426–2457. 10.1257/aer.20110236
- Katz, Lawrence F., and David H. Autor, "Changes in the Wage Structure and Earnings Inequality" (pp. 1463–1555), in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 3A (Amsterdam: Elsevier, 1999). 10.1016/S1573-4463(99)03007-2
- Malamud, Ofer, and Abigail Wozniak, "The Impact of College Education on Geographic Mobility," *Journal of Human Resources* 47:4 (2012), 913–950. 10.1353/jhr.2012.0031
- Mian, Atif, Kamallesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics* 128:4 (2013), 1687–1726. 10.1093/qje/qjt020
- Monte, Ferdinando, Esteban Rossi-Hansberg, and Steve Redding, "Commuting, Migration and Local Employment Elasticities," *American Economic Review* 108:2 (2018). 10.1257/aer.20151507
- Moretti, Enrico, "Local Multipliers," *American Economic Review* 100:2 (2010), 373–377. 10.1257/aer.100.2.373

- “Real Wage Inequality,” *American Economic Journal: Applied Economics* 5:1 (2013), 65–103. 10.1257/app.5.1.65
- Notowidigdo, Matthew J., “The Incidence of Local Labor Demand Shocks,” *Journal of Labor Economics* 38:3 (2020). 10.1086/706048
- Saiz, Albert, “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics* 125:3 (2010), 1253–1296. 10.1162/qjec.2010.125.3.1253
- Schwellnus, Cyrille, Andreas Kappeler, and P. Pionnier, “The Decoupling of Median Wages from Productivity in OECD Countries,” *International Productivity Monitor* (2017), 44–60.
- Sinai, Todd, and Nicholas S. Souleles, “Owner-Occupied Housing as a Hedge against Rent Risk,” *Quarterly Journal of Economics* 120:2 (2005), 763–789.
- Solow, Robert M., “A Contribution to the Theory of Economic Growth,” *Quarterly Journal of Economics* 70:1 (1956), 65–94. 10.2307/1884513
- Stansbury, Anna M., and Lawrence H. Summers, “Productivity and Pay: Is the Link Broken?” NBER working paper 2165 (2017).
- Ströbel, Johannes, and Joseph Vavra, “House Prices, Local Demand, and Retail Prices,” *Journal of Political Economy* 127:3 (2019).
- United Nations, “United Nations Comtrade Database” (2003).
- U.S. Bureau of Labor Statistics, “Relative Importance of Components in the Consumer Price Indexes: U.S.” (Washington, DC: U.S. Bureau of Labor Statistics, 2000).
- Wozniak, Abigail, “Are College Graduates More Responsive to Distant Labor Market Opportunities?” *Journal of Human Resources* 45:4 (2010), 944–970. 10.1353/jhr.2010.0030