Cognitive Hubs and Spatial Redistribution*

Esteban Rossi-Hansberg
Princeton University

Pierre-Daniel Sarte
Federal Reserve Bank of Richmond

Felipe Schwartzman
Federal Reserve Bank of Richmond

February 26, 2021

Abstract

In the U.S., cognitive non-routine (CNR) occupations are disproportionately represented in large cities as well as some smaller cities specializing in CNR intensive industries. To study the allocation of workers across cities, we propose and quantify a spatial equilibrium model with multiple industries employing CNR and non-CNR occupations. Productivity is city-industry-occupation specific and, as we estimate, partly determined by externalities that depend on local occupation shares and total employment. Heterogeneous preferences and these externalities imply inefficient equilibrium allocations. An optimal policy that benefits workers equally incentivizes (i) the formation of cognitive hubs in the largest cities, (ii) higher overall activity and employment in smaller cities, and (iii) increased industrial specialization in both the largest and smallest cities and increases diversification in medium sized cities.

*erossi@princeton.edu, pierre.sarte@rich.frb.org, felipe.schwartzman@rich.frb.org. We thank Mike Finnegan, Xiang Zhang, Daniel Ober-Reynolds, and Jackson Evert for excellent research assistance. We also thank David Albouy, Arnaud Costinot, Morris Davis, Jonathan Dingel, Cecile Gaubert, and Edward Glaeser for useful comments and discussions, as well as participants in various seminars and workshops. The views expressed here are those of the authors and do not reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

1 Introduction

"*Most of what we know we learn from other people (...) most of it we get for free.*"
Robert E. Lucas Jr.[1]

Workers capable of doing the complicated cognitive non-routine tasks required in a modern economy are scarce. Acquiring the expertise to work as a doctor, manager, lawyer, computer scientist or researcher requires many years of schooling, sustained effort, and individual ability. These workers are a valuable input and so their allocation across industries and locations is important for overall efficiency and welfare in an economy. The marginal productivity of a worker depends on the local productivity of the industry where she works, as well as on the set of workers in the same city. Larger cities with a large fraction of workers in cognitive non-routine (CNR) occupations offer learning and collaboration opportunities that enhance the productivity of other workers. However, the abundance of CNR workers also lowers their marginal product, particularly in industries that are less intensive in these occupations. The equilibrium interaction of these forces determines the spatial polarization of workers and, relatedly, the spatial specialization of industries. Can the economy allocate scarce CNR workers in a way that improves the lives of all workers? Our aim is to study the allocation of occupations and industries across cities in the U.S. and to characterize the optimal spatial allocation and the policies to implement it.[2]

The need for optimal spatial policy is the direct implication of the presence of urban externalities. Externalities that enhance the productivity of workers in larger cities have been discussed, analyzed, and measured at least since Marshall (1920).[3] It is natural to hypothesize that these production externalities depend on the occupational composition of a city. After all, CNR occupations require more interactions between knowledgeable workers. As we show in detail in the next section, the patterns of occupational polarization and wages across space in the U.S. suggest that this is indeed the case. First, in the absence of technological differences or externalities across locations, decreasing returns to workers in an occupation imply that relative CNR to non-CNR wages should decline with the share of CNR workers. We find a large positive relationship even after controlling for a number of observable worker characteristics.[4] Why are CNR workers then making relatively more in locations where they are abundant? A possible answer is that these locations specialize in industries intensive in these occupations. The evidence, however, suggests that firms in CNR abundant cities are even more intensive in CNR workers than suggested by their industrial make up. What makes demand for these workers so high in these cities? Our take, and an explanation that reconciles these various facts, is that the abundance of CNR workers itself makes them more productive: a local occupation-specific externality. Estimating the strength of these externalities is a central part of our analysis.

The detailed quantitative assessment of optimal spatial policies we propose requires a number of contributions. These fall along four main dimensions.

First, we develop a spatial equilibrium model with multiple industries and occupations as well as occupation-specific externalities. Multiple industries, costly trade, and input-output linkages are all key features of the environment since the demand for different occupations depends on the occupational intensity of the specific industries in each location. The framework also features

[1] Lucas (1988), page 38.

[2] Figure 20 in the Appendix shows CNR shares across U.S. cities.

[3] See Duranton and Puga (2004) for a review of the literature on externalities in cities.

[4] General city amenities that would be equally appealing to both occupations are unlikely to be a deciding factor in attracting the best workers to abundant cities. Indeed, we show that while real wages of CNR workers increase with a city's CNR intensity, real wages of non-CNR workers do not. Therefore, amenities alone are not driving the pattern of wages across space. For an alternative view, see Couture et al. (2018).

externalities that are allowed to depend on the share of workers in CNR occupations and the total workforce in the city. Finally, it includes heterogeneous preferences for locations that act as a form of migration costs.

Second to arrive at an optimal spatial policy requires that we derive efficient allocations in this setup. We choose to study the efficient allocation that benefit both occupations equally. Implementing this allocation requires particular transfers between locations and occupations, which we characterize.

Third, the details of optimal transfers require that we quantify the model and estimate the parameters that determine the endogenous component of city-industry-occupation specific productivity. Thus, we first recover productivity across locations, industries, and occupation such that the equilibrium of our model matches observed data. We then parameterize the relationship between productivity and the occupational composition and size of cities and estimate this equation using an instrumental variables approach. As proposed in the empirical literature (e.g. Card (2001) and Moretti (2004a)), we use past migration flows of particular immigrant groups and the location of land-grant colleges as instruments. We also present results using model-implied instrumental variables. Our strategy yields a robust set of results comparable to the existing literature though, as an improvement, estimated here directly from productivity measures recovered from the general equilibrium framework we lay out. Our findings imply that the productivity of CNR and non-CNR workers depends similarly on city size. In addition, the productivity of CNR workers depends strongly and significantly on the share of CNR workers. We find less evidence that the productivity of non-CNR workers depends on the composition of occupations.

Finally, informed by these findings, we compute the optimal allocation and discuss its implementation using particular policy tools. We also highlight important quantitative aspects of this allocation through various counterfactual exercises.

Our findings propose a new approach to spatial policy. They indicate that the spatial allocation of workers and industries may be improved by reducing the size of large CNR intensive cities while, at the same time, increasing their fraction of CNR workers. These "cognitive hubs" take advantage of scarce CNR workers in the economy by clustering them to maximize externalities. We find that in equilibrium, the social value of CNR workers is 79% larger than their private value. However, the industrial make up of cities, as well as their location, does impose limits on the creation of cognitive hubs. Some large cities, such as Miami or Las Vegas, remain non-CNR abundant since they are particularly productive in industries where CNR workers are employed less intensively. Cognitive hubs end up scattered geographically around the country to minimize transport cost with the cities with which they trade the most.

In order to increase the share of CNR workers while alleviating congestion in larger cities, the optimal policy prescribes a re-allocation of non-CNR workers to smaller cities with lower CNR shares. The end result is that under the optimal policy, the smallest cities grow in size by playing to their strengths and expanding industries in which a large share of their employment already resides. The corresponding growth of smaller cities also makes it possible for them to sustain more employment in non-tradable industries such as retail, accommodation, and other services. Hence, contrary to some previous literature and much of the public discourse, the economics of the problem suggest that, with the appropriate transfers, small industrial cities in the U.S. should attract non-CNR workers and not try to become the next San Jose. The concentration of CNR workers in a few "cognitive hubs" should be encouraged, not scorned. Everyone can benefit from using CNR workers in the most productive way possible.

Naturally, implementing the optimal allocation requires a number of transfers and taxes that depend on the location and occupation of an agent. As Fajgelbaum and Gaubert (2018) discuss, spatial efficiency requires a flat wage tax on all individuals to correct for the differences in the marginal utility of consumption generated by heterogeneous preferences for location. In addition to this tax, in both our frameworks and theirs, implementing the optimal allocation requires a set of transfers. These transfers insure that non-CNR workers benefit equally from the optimal policy so that occupational inequality is mitigated despite the creation of cognitive hubs. In our analysis, the base transfers to non-CNR workers amount to $16,872 (in 2015 dollars) while CNR workers, who earn substantially more, end up paying a base transfer of $15,255. One interpretation of this base transfer is that of a "universal basic income" paid to all non-CNR workers. CNR workers still need to be incentivized to move to CNR intensive cities, and non-CNR workers to move to non-CNR intensive cities. Therefore, occupation and city-specific transfers are positively correlated with city size for CNR workers and negatively correlated with city size and the CNR share for non-CNR workers. Their exact value depends on the particular location and industrial composition of each city. Ultimately, the policy amounts to a subsidy to non-CNR workers to move to smaller cities with low CNR shares, and incentives to CNR workers to form even more intensive "cognitive hubs" in today's largest cities.

Perhaps surprisingly, a comparison of the current spatial equilibrium to that in 1980 reveals that the spatial allocation of workers has approached that implied by the optimal policy (with current fundamentals). Specifically, since the 1980s, CNR workers have become not only more abundant nationally but also increasingly concentrated in CNR intensive hubs, many of which are large cities. This formation of cognitive hubs has taken place in parallel with a well documented increase in wage inequality across space and occupations. Our quantitative framework implies that both processes were linked through local occupation-specific externalities. Our analysis indicates that absent these spillovers, the spatial polarization of workers would have been greatly mitigated, and the welfare gains received by CNR workers smaller than those of non-CNR workers since CNR workers became more abundant nationally over that time period.

The analysis also makes clear that not all forces pushing towards the spatial polarization of workers are necessarily welfare-enhancing. This is the case, for instance, of housing regulations captured here through changes in the productivity of the real estate sector. The cost of these regulations has been emphasized by, among others, Glaeser and Gyourko (2018), Herkenhoff et al. (2018) and Hsieh and Moretti (2019). Relatively low real-estate productivity growth in CNR-intensive cities since 1980 has increased housing prices and led to more polarized CNR hubs. These changes have thus brought the spatial distribution of occupations closer to that resembling the optimal allocation. In this case, however, since this more spatially polarized distribution of workers resulted from reductions in measured real estate productivity in larger cities, "the cognitive hubs" led to declines in welfare.

Relationship to the Literature A substantial literature has pointed to increasing spatial concentration of skilled workers (Berry and Glaeser (2005), Diamond (2016), and Giannone (2017)), as well as increasing wage inequality across space and within cities (Baum-Snow and Pavan (2013), and Autor (2019)), with the skill premium increasing the most in large cities. Our paper speaks to the optimal policy reaction to those trends.

We focus on production externalities as a key driving force behind those spatial patterns. The estimation of those externalities is a central theme in urban and spatial economics. A robust finding is the existence of a relationship between city size and productivity (see Melo et al. (2009) for a meta analysis). While we allow for such agglomeration externalities, our main focus is on externalities tied to the occupational composition of the city. This is compatible with empirical evidence by Ellison et al. (2010) that industries with similar occupational make-up tend to be spatially proximate. Given the correlation between occupational types and skill levels, our findings of strong spillovers stemming from the occupational composition of cities mirrors findings by Moretti (2004a; 2004b) regarding the local external effects of human capital.

There has been ample research on the extent of spatial misallocation in the U.S. economy and the degree to which it corresponds to heterogeneity in taxation policy (or its local incidence), zoning laws, or other unspecified sources of distortions. Examples of papers in that vein are Albouy (2009), Desmet and Rossi-Hansberg (2013), Ossa (2015), Fajgelbaum et al. (2018), Colas and Hutchinson (2017), Hsieh and Moretti (2019) and, most recently Herkenhoff et al. (2018).

Our paper sheds light on place-based policies in that it highlights the optimal endogenous expansion of different industries in different locations. A summary of the related literature can be found in Neumark and Simpson (2015). Rather than evaluating exogenous policies, we derive the optimal allocation in a quantitative spatial model with local externalities. Our derivation of optimal policy thus generalizes that of Fajgelbaum and Gaubert (2018) in an environment with input-output linkages and where trade is differentially costly across industries. Two other recent papers that discuss the optimal distribution of city sizes are Eeckhout and Guner (2015) and Albouy et al. (2019).

We integrate industrial, occupational and spatial heterogeneity in a single coherent framework. Other recent work that has emphasized the joint distribution of industrial and skill composition within the U.S. are Hendricks (2011) and Brinkman (2014). As in Caliendo et al. (2017), we allow for trade costs, thus capturing an explicitly spatial dimension, but add to that framework by also allowing for occupational heterogeneity and local production externalities. Finally, on a more technical note, our paper adds to the rapidly expanding 'quantitative spatial economics' literature that uses general equilibrium models to address issues related to international, regional and urban economics. Redding and Rossi-Hansberg (2017) provide a review of the main ingredients in these models.

The rest of the paper is organized as follows. Section 2 presents stylized facts that constitute prima-facie evidence for the presence of externalities among CNR workers within cities. Section 3 presents our multi-industry spatial model with occupation specific externalities within cities. Section 4 quantifies the model, including our estimation of the externality parameters. It also discusses the role of externalities in the equilibrium allocation. Section 5 presents the optimal allocation as well as the resulting transfers and their implementation. Section 6 provides a decomposition of the impact of fundamental changes in the national CNR employment share and in technology across sectors and cities between 1980 and the recent data. Section 7 concludes. We relegate many of the model's details, additional robustness exercises and counterfactuals to the Appendix.

2 A Motivation for Occupation-Specific Externalities

The main question under consideration is whether there is a role for policy in altering the observed spatial polarization of employment and, if so, what are its features? We now provide some basic facts regarding the joint spatial distribution of wages and employment for workers in different occupations that point to the existence of important occupational externalities. Those facts constitute prima-facie evidence that the optimal policy may in fact involve reinforcing existing patterns, with the appropriate transfers, rather than attenuating them.

We separate workers from 2011 to 2015 in two large occupational groups: those that are intensive in cognitive non-routine (CNR) tasks and the others (non-CNR).[5] We calculate the average residual

[5]Specifically, we follow Jaimovich and Siu (2018), and define CNR occupations to include occupations with SOC-2 classifications 11 to 29 and non-CNR occupations to include the remainder of SOC-2 classifications. Wage and employment data is obtained from the American Community Survey.

Figure 1: Occupational employment share and wage premium

See text for details on definitions of the wage premium and occupation classification. Log wage premium is depicted as a deviation from employment weighted mean. Each observation refers to a CBSA. Marker sizes are proportional to total employment.

wages of workers in each occupation and each city after controlling for observable worker socio-economic characteristics.[6] This classification builds on the observation by Acemoglu and Autor (2011) that one can best understand wage inequality trends through such a task based approach.

Figure 1 shows that across U.S. cities, wages of workers employed in Cognitive Non-Routine (CNR) occupations, relative to those of workers in other (non-CNR) occupations, increase with the corresponding share of CNR workers in total employment. This suggests that differences in relative wages across cities are, to a large degree, driven by differences in relative demand for CNR workers.[7] The size of the scatter-plot markers captures city size. They indicate that large cities appear to also be CNR intensive.

Focusing on CNR workers, the left panel of Figure 2 indeed shows that real wages of CNR workers increase with the intensity of CNR employment across cities. Moreover, some of the high real wage cities include places like San Francisco and New York that on average may provide higher amenities to CNR workers (see Diamond (2016)). In those cities, therefore, labor demand forces are seemingly pronounced enough to more than make up for the labor supply inducing effects of local amenities, such as the variety of retail and entertainment options. If workers differ in their preferences for where to live, the real wages depicted in the left panel of Figure 2 reflect the compensating differential to the marginal CNR worker in a given city.

Differences in the relative demand for CNR workers across cities can arise for several reasons.

[6]We include as control variables education, potential experience, race, gender, English proficiency, number of years in the U.S., marital status, having had a child in the last year, citizenship status, and veteran status.

[7]In particular, suppose that technologies were similar across cities, and that the share of CNR workers were driven by the supply of those workers. Then, with decreasing marginal returns to worker type, increases in the relative supply of CNR workers would lower their relative wages.

First, differences in relative demand for CNR workers may arise from exogenous (or historically determined) differences in industrial composition or regulations. Suppose that the industry make up of a city, n, is the main determinant of its demand for CNR workers relative to other types. Then, its wage bill share for CNR workers would be (approximately) $\sum_j \delta^{CNR,j}\sigma_n^j$, where $\delta^{CNR,j}$ is the national wage bill share of CNR workers in industry j, and σ_n^j is the wage bill of industry j as a share of that city's total wage bill. Figure 3 compares the wage-bill shares of CNR workers implied by the different industrial composition of U.S. cities relative to those observed in the data. The black line is a 45 degree line. The observed wage bill shares differ from, and in fact increase more than one-for-one with, those implied by differences in industrial mix alone, thus ruling out industrial composition as a sole determinant of labor demand across cities.[8]

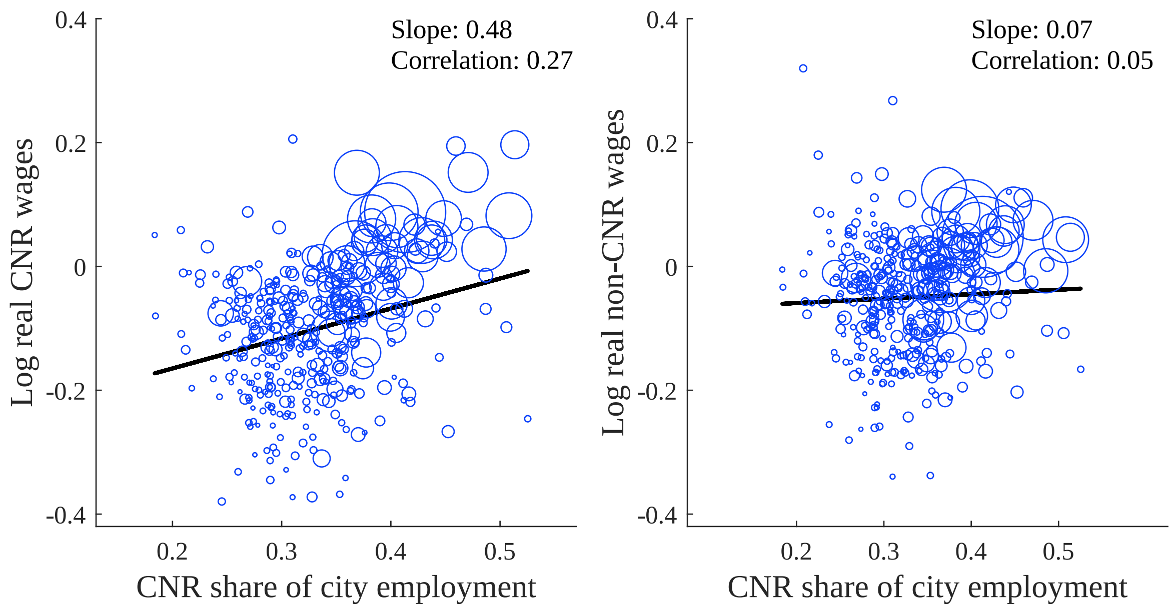

Figure 2: Occupational employment share and real wages

See text for details on definitions of the wage premium and occupation classification. Real wages are calculated using consumption price indices obtained from the model quantification (See Section 4). Log real wages are depicted as a deviation from employment weighted mean. Each observation refers to a CBSA. Marker sizes are proportional to total employment.

Differences in the relative demand for CNR workers across cities can also arise endogenously if more productive workers within occupational types sort themselves into particular cities. Baum-Snow and Pavan (2013) indeed argue that observable worker characteristics are an important determinant of the city size wage premium. However, the fact that relative wages in Figure 1 are computed from residuals after controlling for observable worker characteristics suggests that sorting along these characteristics is not the only driving force underlying that figure. Thus, if sorting is nevertheless part of an explanation driving the positive relationship between the wage premia of CNR workers and the employment share of those workers, it must be taking place along dimensions that are not easily observed. However, assuming that differences in amenities are experienced in similar ways by CNR and non-CNR workers, high productivity non-CNR workers would then also

[8] The figure also rules out the production technology for different industries being well characterized by Cobb-Douglas (i.e. the elasticity of substitution across worker types is likely not equal to 1).

sort themselves into cities with a high share of CNR workers. The right panel of Figure 2 suggests that this is not, in fact, the case.[9]

Finally, differences in relative demand for CNR workers may be explained by endogenous differences in productivity, even when not from sorting, if these differences arise from production externalities that predominantly affect CNR workers. First, to the extent that production externalities also increase with the concentration of CNR workers in a given city, it is then naturally the case that the demand for CNR workers would increase with the share of employment in CNR occupations, as suggested by Figure 1. Second, if production externalities mainly enhance the productivity of CNR workers, then real wages of CNR workers would increase with the share of CNR employment within cities, as in the left panel of Figure 2, but no such effect would be expected among non-CNR workers, as suggested by the right panel of Figure 2. Third, and most importantly, Figure 3 shows that observed wage bill shares of CNR workers increase more than one-for-one with those implied by differences in industrial composition alone. This observation would be expected in an environment where production externalities intensify the implications of industrial mix. Specifically, CNR workers will concentrate, all else equal, in cities whose industrial composition is tilted towards industries intensive in CNR workers. In the presence of production externalities, therefore, this concentration would lead to increases in the productivity of CNR workers. If the elasticity of substitution between worker types is higher than 1, one would then expect higher wage shares for CNR workers in those cities relative to those given by industrial composition alone.

3 A Quantitative Spatial Model with Multiple Industries and Occupations

The economy has N cities and J sectors. We denote a particular city by $n \in \{1, ..., N\}$ (or n') and a particular sector by $j \in \{1, ..., J\}$ (or j'). Individuals are endowed with an occupational type and cannot switch types. There are K occupational types, denoted by $k \in \{1, ..., K\}$ (or k'), with aggregate number of workers L^k per type (total employment in occupation k aggregated across industries and cities).[10] Firms in all cities use multiple types of labor but in potentially different proportions depending on the industry and the city. Aggregate regional land and structures in region n are denoted by H_n. Labor of all types moves freely across regions and sectors, while structures are region-specific. Some sectors are tradable while others are not.

Quantities in the economy may be associated with industries, cities, or occupations. For notational convenience, we denote aggregates across a given dimension by omitting the corresponding index. Thus, for example, L_n^{kj} is the number of workers employed in occupation k in industry j in city n, $L_n^k = \sum_j L_n^{kj}$ represents the number of workers employed in occupation k in city n, $L^k = \sum_n L_n^k$ represents all workers in occupation k, and $L = \sum_k L^k$ is simply total employment.

3.1 Individuals

Workers in each location $n \in \{1, ..., N\}$, are endowed with labor of type $k \in \{1, ..., K\}$, and order consumption baskets according to Cobb-Douglas preferences with shares α^j over their consumption of final domestic goods, C_n^{kj}:

[9]The small relevance of sorting to explain differences in wages across cities has in fact been recently verified in empirical work by Baum-Snow and Pavan (2011) and Roca and Puga (2017).

[10]The two broad occupation groups that we examine in general require different levels of education and skill. We do not model the occupational choice of these workers but could potentially do so, as in Burstein et al. (2019)

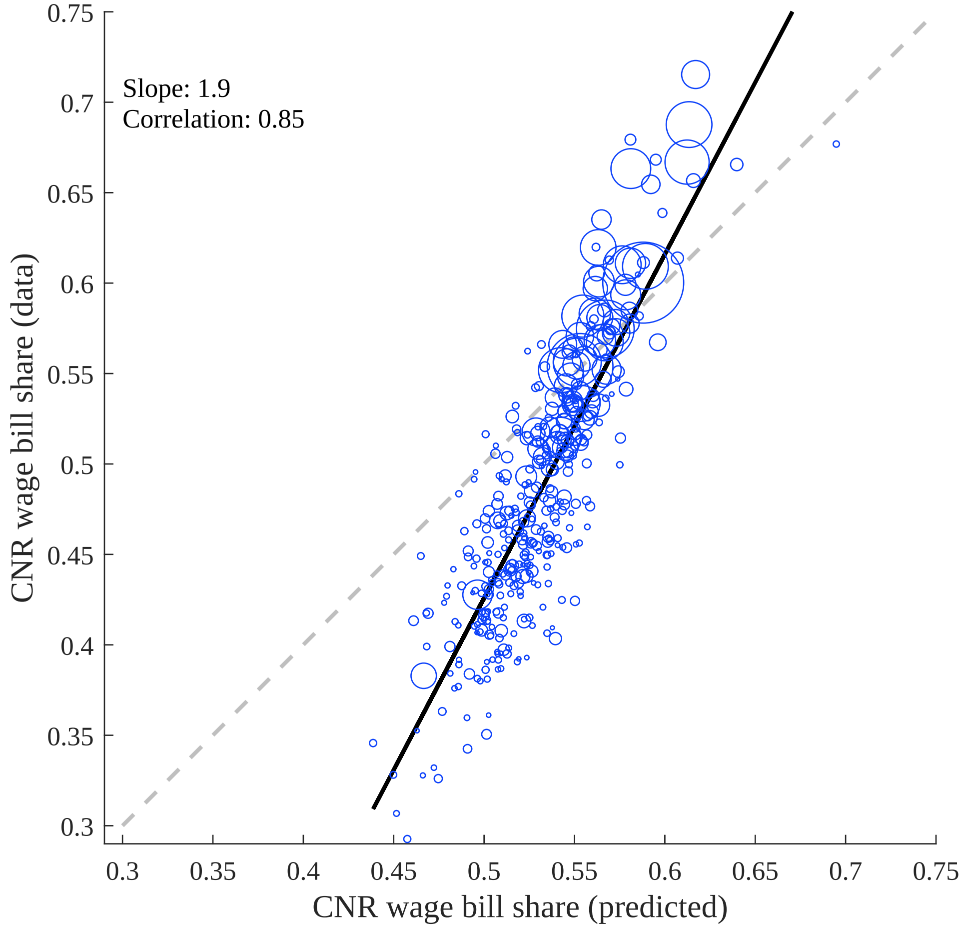

Figure 3: Occupational wage bill share: predicted vs actual
See text for details on definitions of the wages and occupation classification. The predicted wage bill-shares are obtained by assuming that within-industry wage bill shares were equal to national averages (see text for details). Each observation refers to a CBSA. Marker sizes are proportional to total employment.

$$C_n^k = \prod_j \left(C_n^{kj} \right)^{\alpha^j},$$

where C_n^k is a consumption aggregator. Consumption goods consumed in city n are purchased at prices P_n^j in sectors $j \in \{1, ..., J\}$. Utility is homogeneous of degree one, so that $\sum_j \alpha^j = 1$.

Workers supply one unit of labor inelastically. The income of a worker of type k residing in city n is

$$I_n^k = w_n^k + \chi^k, \tag{1}$$

where w_n^k is the wage earned by a worker in occupation k in city n. The term χ^k represents the return per household on a national portfolio of land and structures from all cities,

$$\chi^k = b^k \frac{\sum_{n'} r_{n'} H_{n'}}{L^k},$$

where r_n is the rental rate on land and structures in that city, and b^k denotes the share of the national portfolio accruing to workers of occupational type k. In what follows, we assume that b^k is determined such that different worker types receive a share of the national portfolio proportional to their share of wages in the total wage bill, so

$$b^k = \frac{\sum_n w_n^k L_n^k}{\sum_{k',n'} w_{n'}^{k'} L_{n'}^{k'}}.$$

Agents of a given occupational type differ in how much they value living in different cities. These differences are summarized by a vector $\mathbf{a} = \{a_1, a_2, ..., a_N\}$, with each entry a_n scaling the utility value that an individual receives from living in city n. We associate the elements a_n with the particular way in which different workers experience the amenities of given cities. Conditional on living in city n, the problem of an agent employed in occupation k and characterized by amenity vector $\mathbf{a}$ is

$$v_n^k(\mathbf{a}) \equiv \max_{\left\{C_n^{kj}(\mathbf{a})\right\}_{j=1}^J} a_n A_n^k \prod_j \left(C_n^{kj}(\mathbf{a})\right)^{\alpha^j}, \text{ subject to } \sum_j P_n^j C_n^{kj}(\mathbf{a}) = I_n^k,$$

where A_n^k denotes an exogenous component of city-specific utility common to all individuals of type k living in city n. All workers in a given occupation living in a given city will choose the same consumption basket. It follows that $C_n^{kj}(\mathbf{a}) = C_n^{kj}$ for all $\mathbf{a}$.

Agents move freely across cities. The value of locating in a particular city n for an individual employed in occupation k, with idiosyncratic preference vector $\mathbf{a}$, is

$$v_n^k(\mathbf{a}) = \frac{a_n A_n^k I_n^k}{P_n} = a_n A_n^k C_n^k.$$

In equilibrium, workers move to the location where they receive the highest utility so that

$$v^k(\mathbf{a}) = \max_n v_n^k(\mathbf{a}),$$

where $v^k(\mathbf{a})$ now denotes the equilibrium utility of an individual in occupation k with amenity vector $\mathbf{a}$. We assume that a_n is drawn from a Fréchet distribution independently across cities. We denote by Ψ the joint *cdf* for the elements of $\mathbf{a}$ across workers in occupation k, so that

$$\Psi(\mathbf{a}) = \exp\left\{-\sum_n (a_n)^{-\nu}\right\},$$

where the shape parameter ν reflects the extent of preference heterogeneity across workers employed in occupation k. Higher values of ν imply less heterogeneity, with all workers ordering cities in the same way when $\nu \to \infty$. The assumption of a Fréchet distribution for idiosyncratic amenity parameters implies closed form expressions for the fraction of workers in each city:

$$L_n^k = \Pr\left(v_n^k(\mathbf{a}) > \max_{n' \neq n} v_{n'}^k(\mathbf{a})\right) = \frac{\left(A_n^k C_n^k\right)^\nu}{\sum_{n'}\left(A_{n'}^k C_{n'}^k\right)^\nu} L^k. \tag{2}$$

3.2 Firms

There are two types of firms: those producing intermediate goods and those producing final goods. There is a continuum of varieties of intermediate goods which are aggregated into a finite number of final goods corresponding to J sectors. Varieties of intermediate goods are characterized by the sector in which they are produced, and by a vector of city-specific productivity parameters, $\mathbf{z} = \{z_1, z_2, ..., z_N\}$, with each element z_n scaling the productivity of firms in city n producing that variety.

Final goods are sold in the city where they are produced. Varieties of intermediate goods are traded across cities. Because of transportation costs, the price earned by intermediate goods producers need not be the same as the price paid by final goods producers. Intermediate goods producers operating in city n, sector j, producing a variety indexed by $\mathbf{z}$, produce a quantity, $q_n^j(\mathbf{z})$, for which they earn a price $p_n^j(\mathbf{z})$. Final goods producers operating in city n, sector j, purchase a quantity $Q_n^j(\mathbf{z})$ of the variety of intermediate goods indexed by $\mathbf{z}$.

3.2.1 Intermediate Goods

Idiosyncratic productivity draws, $\mathbf{z}$, arise from a Fréchet distribution with shape parameter θ. Draws are independent across goods, sectors, and regions. Specifically, if we let Φ be the joint *cdf* of variety-specific productivity parameters across firms in industry j, then

$$\Phi(\mathbf{z}) = \exp\left\{-\sum_n (z_n)^{-\theta}\right\}.$$

Production of intermediate goods a variety indexed by $\mathbf{z}$, in city n, and industry j, takes place using the technology,

$$q_n^j(\mathbf{z}) = z_n \left[H_n^j(\mathbf{z})^{\beta_n^j} \left[\sum_k \left(\lambda_n^{kj} L_n^{kj}(\mathbf{z}) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}(1-\beta_n^j)} \right]^{\gamma_n^j} \prod_{j'} M_n^{j'j}(\mathbf{z})^{\gamma_n^{j'j}} \tag{3}$$

where $\gamma_n^{j'j} \geq 0$ is the share of sector j input expenditures spent on materials from sector j' in city n, $\gamma_n^j \geq 0$ is the share of value added in gross output in sector j, and β_n^j is the share of land and structures in value added in that sector. The production function is constant returns to scale, $\sum_{j'=1}^J \gamma_n^{j'j} = 1 - \gamma_n^j$. The variable λ_n^{kj} denotes a labor augmenting productivity component that is city, industry, and occupation specific. We denote by $H_n^j(\mathbf{z})$ the quantity of structures used by a firm producing a variety $\mathbf{z}$ in industry j operating in city n, by $M_n^{j'j}(\mathbf{z})$ the quantity of material goods this firm uses from sector j', and by $L_n^{kj}(\mathbf{z})$ the workers of type k it employs.

Importantly, given the evidence presented in Section 2, we allow λ_n^{kj} to reflect externalities that depend on the composition of the labor force. In particular, we let

$$\lambda_n^{kj} = \lambda_n^{kj}(\mathbf{L_n}),$$

where $\mathbf{L}_n = \{L_n^1, ..., L_n^K\}$ summarizes the occupational make up of the labor force in city n.

3.2.2 Final Goods

A final goods firm operating in industry j in city n produces the quantity Q_n^j according to the technology,

$$Q_n^j = \left[\int \left[\sum_{n'} Q_{nn'}^j(\mathbf{z}) \right]^{\frac{\eta-1}{\eta}} d\Phi(\mathbf{z}) \right]^{\frac{\eta}{\eta-1}},$$

where $Q^j_{nn'}(\mathbf{z})$ represents its use of intermediate goods of variety $\mathbf{z}$ produced in city n'.

One unit of any intermediate good in sector j shipped from region n' to region n requires producing $\kappa^j_{nn'} \geq 1$ units in the origin n'. Therefore, producers of final goods in each sector solve

$$\max_{Q^j_{nn'}(\mathbf{z})} P^j_n Q^j_n - \sum_{n'} \int \kappa^j_{nn'} p^j_{n'}(\mathbf{z}) Q^j_{nn'}(\mathbf{z})\, d\Phi(\mathbf{z}),$$

subject to $Q^j_{nn'}(\mathbf{z}) \geq 0$, where P^j_n is the price of the final good in sector j, city n. Intermediate goods in non-tradable sectors cannot be shipped between cities.

Final goods firms purchase intermediate goods from the location in which the acquisition cost, including transportation costs, is the least. Denote by X^j_n the total expenditures on final goods j by city n, which must equal of the value of final goods in that sector, $X^j_n = P^j_n Q^j_n$. Because of zero profits in the final goods sector, total expenditures on intermediate goods in a given sector are then also equal to the cost of inputs used in that sector. Following the usual Eaton and Kortum (2002) derivations, given a final good j produced in city n, the share of intermediate inputs imported from location n' is

$$\pi^j_{nn'} = \frac{\left[\kappa^j_{nn'} x^j_{n'}\right]^{-\theta}}{\sum_{n''=1}^{N} \left[\kappa^j_{nn''} x^j_{n''}\right]^{-\theta}},$$

where

$$x^j_n = \left\{ \left(\frac{r_n^{\beta^j_n}}{\beta^j_n} \right) \left[\frac{1}{1-\beta^j_n} \sum_k \left(\frac{w^k_n}{\lambda^{kj}_n} \right)^{1-\epsilon} \right]^{\frac{1-\beta^j_n}{1-\epsilon}} \right\}^{\gamma^j_n} \prod_{j'=1}^{J} \left(\frac{P^{j'}_n}{\gamma^{j'j}_n} \right)^{\gamma^{j'j}_n} \tag{4}$$

is a cost index associated with the production of varieties in sector j, city n. In quantifying the model, we also allow for two non-tradable sectors for which $\pi^j_{nn} = 1$.

3.3 Market Clearing Conditions

Within each city n, the number of workers employed in occupation k must equal the number of those workers who choose to live in that city. Put alternatively,

$$\sum_j \int L^{kj}_n(\mathbf{z}) d\Phi(\mathbf{z}) = \int \zeta^k_n(\mathbf{a}) d\Psi(\mathbf{a}),\ \forall\ n = 1, ..., N,\ k = 1, ..., K. \tag{5}$$

where $\zeta^k_n(\mathbf{a}) \in \{0, 1\}$ denotes the location choice of households as a function of their type. Market clearing for land and structures in each region imply that

$$\sum_j \int H^j_n(\mathbf{z})\, d\Phi(\mathbf{z}) = H_n,\ n = 1, ..., N. \tag{6}$$

Final goods market clearing implies that

$$\sum_k L^k_n C^{kj}_n + \sum_{j'} \int M^{jj'}_n(\mathbf{z})\, d\Phi(\mathbf{z}) = Q^j_n. \tag{7}$$

Finally, intermediate goods market clearing requires that

$$q_n^j(\mathbf{z}) = \sum_{n'} \kappa_{n'n}^j Q_{n'n}^j(\mathbf{z}). \tag{8}$$

4 Quantifying the Model

In the model we have laid out, any quantitative statement about efficient allocations will necessarily depend on the parameterization of occupation-specific externalities. Estimating these externalities in turn requires that we take the model to the data and recover productivity in different occupations, sectors, and cities such that the equilibrium of the model matches these data. Observations used in the model inversion are thus matched by construction. However, other predictions of the model with respect to recovered productivity or tradable goods prices do not have easily observable counterparts. In fact, we show that the properties of productivity and prices delivered by the model are comparable to those found in different recent studies, effectively tying these studies within a single general equilibrium framework.

The model is mapped into 22 industries and the two large occupational groups (cognitive non-routine and others) emphasized in Section 2. Of the 22 industries, two are non-tradable, meaning that all local output is also used locally. The two non-tradable sectors include real-estate services, which is the single user of land in each city, and a composite sector comprising retail, construction, and utilities. Tradable industries include 10 manufacturing sectors and 10 service sectors. In quantifying the model, we focus on the period 2011 to 2015.

The set of parameters needed to quantify our framework fall into two broad types: i) parameters that are constant across cities (but may vary across occupations and/or industries) and ii) parameters that vary at a more granular level and require using all of the model's equations (i.e. by way of model inversion) to match data that vary across cities, industries, and occupations.

To obtain an initial calibration for the share parameters γ_n^j, $\gamma_n^{jj'}$ and α^j, we use an average of the 2011 to 2015 BEA Use Tables, each adjusted by the same year's total gross output, where we assume that tradable sectors have $\gamma_n^j = \gamma^j$ that are constant across cities and similarly for $\gamma_n^{jj'}$'s.[11]

We adopt the convention that all land and structures are managed by firms in the real estate sector that then sell their services to other sectors. Accordingly, for all sectors other than real estate, we reassign the gross operating surplus remaining after deducting equipment investment to purchases from the real estate sector. These surpluses are in turn added to the gross operating surplus of real estate.[12] This convention implies that the share of land and structures, β_n^j, in the production of all sectors other than real estate is equal to zero. Observe also that the share, β_n^j, helps determine the supply elasticity of real estate services which differs across cities.

We set θ, the Fréchet parameter governing trade elasticities, to 10 or well within the range of estimates of trade elasticities in the literature, between 3 to 17 (see Footnote 44 in Caliendo and Parro (2015), as well as Head and Mayer (2014), section 4.2 for comprehensive summaries of estimates). While estimates of θ have been carried out at various levels of disaggregation, these can vary somewhat widely for a given sector or commodity across studies.[13] For our purposes, this

[11] Since the model does not allow for foreign trade, we adjust the Use Table by deducting purchases from international producers from the input purchases and, for accounting consistency, from the definition of gross output for the sector.

[12] One can verify that those reassignments do not affect aggregate operational surplus (net of equipment investment), aggregate labor compensation, and aggregate value added (net of equipment investment).

[13] For example, while Caliendo and Parro (2015) estimate an elasticity of 7.99 for Basic Metals and 4.75 for Chemicals, Feenstra et al. (2018) estimate a elasticities of, respectively, 1.16 and 1.46 for those two categories.

uncertainty is further compounded by the fact that trade elasticities that are relevant for trade between countries may not be appropriate for trade between regions or cities.

As mentioned, we assume that two of the sectors ("real estate," as well "retail, construction, and utilities") are non-tradable, so that their transportation costs are treated as infinite. For the tradable sectors, we follow Anderson et al. (2014) and assume that trade costs increase with distance. Specifically, in order to ship one unit of good to city n, $\kappa^j_{nn'} = (d_{nn'})^{t^j}$ units of the good need to be produced in city n', with $d_{nn'}$ the distance between city n and city n' in miles.[14] The parameter t^j is industry specific. For commodities, we directly estimate t^j from the Commodity Flow Survey synthetic microdata using standard gravity regressions based on model trade-shares. In the tradable services, we use the values obtained by Anderson et al. (2014) using Canadian data.

Ciccone and Peri (2005) summarize estimates for the elasticity of substitution between skilled and unskilled labor in the literature as ranging between 1.36 and 2. Card (2001) estimates the elasticity of substitution between occupations to be closer to 10. We adopt $\epsilon = 2$ as a benchmark. Finally, we set ν so that the average elasticity of employment with respect to real wages in our model matches the estimate of 1.36 as in Table A.11, column 4, of Fajgelbaum et al. (2018). This implies $\nu = 2.02$.[15]

Given the parameters above, we use data on wages by occupation and location (w^k_n), as well as data on employment by occupation, industry and location (L^{kj}_n), to obtain equilibrium values of productivity and amenities, λ^{kj}_n and A^k_n respectively. Data pertaining to w^K_n, and $\frac{L^{kj}_n}{\sum_{k'} L^{k'j}_n}$ is available from the American Community Survey (ACS). The ACS also allows us to adjust wages for individual characteristics so that our data captures city wage premia for each occupation. The Census provides measures of total employment, $\sum_{k'} L^{k'j}_n$, from the County Business Patterns (CBP) that better match BEA industry-level counts. We combine total employment from the CBP with ACS data on employment shares to obtain L^{kj}_n. The exact procedure that yields λ^{kj}_n and A^k_n by way of model inversion is described in detail in Appendix B.

Table 1 compares the relationships between wages, employment, and employment composition across different cities highlighted in previous work relative to the data used in our model inversion. The first three rows of the table show regression coefficients of log wages for CNR workers, non-CNR workers, and the CNR wage premium, on different measures of city employment and employment composition. The subsequent rows show similar regression coefficients obtained in previous literature. The data we use implies relationships that are consistent with those in other work. In particular, all wages increase with city size, more so for skilled workers. A similar relationship holds for wages and city composition where proportionally more skilled cities exhibit higher wages for all workers, more so for skilled workers.[16]

4.1 Model Validation

We now show that our model-consistent TFP measures and tradable prices compare favorably with previous empirical work, but within a single general equilibrium framework that can also be used

[14] We assume that within city distance is equal to 20 miles.

[15] Here, ν is somewhat larger than the value obatained by Fajgelbaum et al. (2018). This reflects the fact that ν is the elasticity of labor supply with respect to consumption rather than wages. Because Fajgelbaum et al. (2018) abstract from non-wage income, they estimate values for the analogous parameter in their model of between 0.75 and 2.25 depending on identification assumptions.

[16] An exception is Moretti (2004a) who finds no statistically significant differences in the way that wages of college educated workers and non-college educated workers vary with employment composition across cities. Our findings, however, rely on a more recent time period where other work has found an increasingly pronounced relationship between skill and city size (see Baum-Snow and Pavan (2013)).

to guide optimal policy. There exists a large literature that has estimated and studied the role of agglomeration externalities. Much of this work has relied on a production function approach using measures of output and factor inputs to estimate Total Factor Productivity (TFP), or using labor productivity more directly, in exploring how productivity depends on the scale of city employment or its skill composition. There is also a literature that has sought to understand how tradable goods prices vary with city size. We add to those literatures by combining their results with a framework that can then be used to provide a quantitative assessment of optimal spatial policy.

4.1.1 Tradable Goods Prices and City Size

Recent work by Handbury and Weinstein (2014), using Nielsen home-scanned data on tradable goods bought in grocery stores, highlights that tradable consumer prices decrease with city size. Prior to that study, the consensus view, based on more aggregated prices, was that such prices

Table 1: Wages, Employment, and City Composition

Dependent Variable	$\ln(L_n)$	$\ln\left(\frac{L_n^{\text{CNR}}}{L_n^{\text{nCNR}}}\right)$	$\frac{L_n^{\text{CNR}}}{L_n}$
$\ln\left(w_n^{\text{CNR}}\right)$	0.059 (0.004)	0.338 (0.021)	1.499 (0.089)
$\ln\left(w_n^{\text{nCNR}}\right)$	0.050 (0.002)	0.223 (0.017)	0.993 (0.072)
$\ln\left(\frac{w_n^{\text{CNR}}}{w_n^{\text{nCNR}}}\right)$	0.022 (0.001)	0.115 (0.008)	0.506 (0.036)
Moretti HS[1]	—	—	0.85 (0.06)
Moretti Some College	—	—	0.86 (0.06)
Moretti College +	—	—	0.74 (0.06)
Roca & Puga wage log wage constant[2]	0.0455 (0.0080)	—	—
Diamond log college wage[3]	—	0.26 (0.11)	—
Diamond log non-college wage[4]	—	0.18 (0.01)	—
Baum-Snow et al. log wage, 2005-2007[5]	0.065 (< 0.01)	—	—
Baum-Snow et al. log wage ratio[6]	0.029 (< 0.003)	—	—

1. Moretti (2004a) "Estimate the social return to higher education: evidence from longitudinal and repeated cross-sectional data", Table 5.
2. Roca and Puga (2017) "Learning by Working in Big Cities", Table 1.
3. Diamond (2016) "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000", Figure 4.
4. Diamond (2016), Figure 3
5. Baum-Snow et al. (2018) "Why Has Urban Inequality Increased?", Table 1. Standard error reported as less than 0.01.
6. Baum-Snow et al. (2018), Table 2. Standard error reported as less than 0.003.

instead increased with city size.[17] Given the detailed nature of Nielsen home-scanned prices, Handbury and Weinstein (2014) are able to control for product buyer and retailer heterogeneity in a way that is not easily achieved with more aggregate prices. Allowing for those controls reduces the elasticity of tradable goods bought in grocery stores with respect to city size to zero. When Handbury and Weinstein (2014) further adjust local price indices to reflect differences in the number of varieties of goods available in different cities, they find that the price of tradable goods bought in grocery stores actually decreases with city size, with an elasticity equal to -0.011 (Table 6). In addition, when calculating this elasticity after purging the effect of local rents on retail costs, they obtain -0.017 (Table 9 in Working Paper version). In the presence of trade costs, such a declining relationship indeed emerges when larger cities are generally more productive. It can hold for certain categories of goods even if the local scarcity of land means that the price of real estate services is higher, thus driving up the general price level, in larger cities.

Table 2 below summarizes the relationship between prices obtained in our model inversion and city size. Similar to Handbury and Weinstein (2014), our general equilibrium framework reveals a decreasing relationship between prices and city size and, in fact, across all tradable sectors with an average elasticity of -0.012. In the Food and Beverage sector, our model inversion reveals an elasticity of -0.010, virtually identical to that Handbury and Weinstein (2014) for grocery products. Remarkably, our finding arises without direct observation of prices. Instead they follow from supply and demand relationships within a structural trade model where cities produce different goods and where trade across regions is costly. When informed by the data described above, our model then implies that large cities are generally more productive thus yielding smaller prices for tradable goods.

4.1.2 Amenities

We now turn to the occupation-specific amenities implied by the model inversion, A_n^k. The relationship between relative amenities for CNR and non-CNR workers against the size and composition of cities is depicted in Figure 4. Our findings conform to Diamond (2016) in that cities with more CNR workers are also relatively more amenable to those same workers. At the same time, larger cities are relatively more amenable to CNR workers helping account for the concentration of CNR workers in large cities.

Diamond (2016) provides evidence for a causal impact of local population composition on amenities. In Appendix E, we show the effect of filtering out the component of amenities that is endogenous to the local labor composition.[18] While suppressing those endogenous effects eliminates the positive relationship between the CNR share and relative amenities, the relationship between relative residual amenities and city size becomes stronger. Intuitively, given the estimates in Diamond (2016), large non-CNR populations generate larger congestion effects on CNR workers than on non-CNR workers. The bottom line, therefore, is that our findings below regarding the optimality of concentrating CNR workers, computed without endogenous amenities, are if anything conservative.

4.1.3 Total Factor Productivity

A substantive literature in urban economics has addressed the relationship between productivity and city size (i.e. "agglomeration economies"), as well as that between productivity and employment composition. Baseline estimates of real Total Factor Productivity (TFP) typically rely

[17] In principle, given that rents generally increase with city-size, tradable consumer prices might indeed follow the same pattern to the degree that they are partially influenced by local rents as an input cost.

[18] Here, we use the parameterization that Fajgelbaum and Gaubert (2018) obtain based on the estimates by Diamond (2016).

Table 2: Elasticities of Final Goods Prices, P_n^j, w.r.t. L_n

Sector	Elasticity
Food and Beverage	**−0.010**
Textiles	−0.022
Wood, Paper, and Printing	−0.011
Oil, Chemicals, and Nonmetallic Minerals	−0.016
Metals	−0.014
Machinery	−0.005
Computer and Electronic	−0.013
Electrical Equipment	−0.003
Motor Vehicles (Air, Cars, and Rail)	−0.009
Furniture and Fixtures	−0.009
Miscellaneous Manufacturing	−0.013
Wholesale Trade	−0.007
Transportation and Storage	−0.007
Professional and Business Services	−0.018
Other	−0.011
Communication	−0.003
Finance and Insurance	−0.011
Education	−0.024
Health	−0.030
Accommodation	−0.012
Real Estate	0.131
Retail, Construction and Utilities	0.047
Average	−0.003
Tradable Average	−0.012
Manufacturing Average	−0.011
Tradable Services Average	−0.014

Coefficients from univariate OLS regression of final-goods prices ($\ln P_n^j$) on $\ln L_n$.

on Cobb-Douglas production functions that allow for different types of labor to enter separately. Within the context of our model, we follow Caliendo et al. (2017) and express measured TFP as

$$\ln TFP_n^j = \ln \frac{\sum_{n'} \pi_{n'n} X_{n'}^j}{P_n^j} - \gamma_n^j \beta_n^j \ln H_n^j - \gamma_n^j (1 - \beta_n^j) \sum_k \delta^{kj} \ln L_n^{kj} - \sum_{j'} \gamma_n^{j'j} \ln M_n^{j'j}, \quad (9)$$

where δ^{kj} is the share of occupation k wages in sector j's wage bill. In the model we have laid out, and up to a first-order approximation (abstracting from selection effects induced by trade), we have that for tradable sectors,

$$\ln TFP_n^j \simeq \sum_k \delta^{kj} \gamma^j \ln \lambda_n^{kj},$$

where $\left(\lambda_n^{kj}\right)^{\gamma^j}$ may thus be interpreted as the component of TFP in sector j and city n associated with occupation k.[19] In the remainder of the paper, we let $T_n^{kj} = \left(\lambda_n^{kj}\right)^{\gamma^j}$.

Table 3 below shows the city in which TFP recovered from the model is highest by industry. The results largely conform to intuition. Productivity in Computers and Electronic Equipment is

[19] See Appendix B for the details of this derivation. For non-tradable sectors, the city-specific share parameters make it challenging to compare this term across-cities. Furthermore, our data does not allow us to separate the productivity of the real-estate sector from the stock of housing.

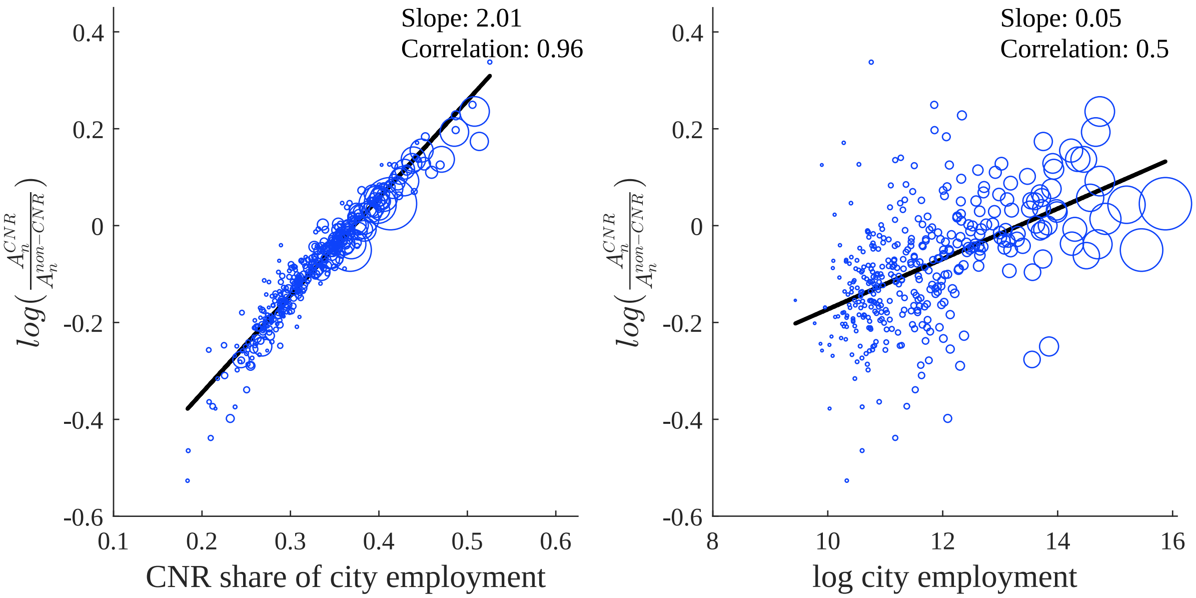

Figure 4: Relative amenities and city size and composition

Ratio of occupational-specific amenity parameters for each city obtained from the model quantification against employment share of CNR workers and log total employment. Each observation refers to a CBSA. Marker sizes are proportional to total city employment.

highest in San Jose, CA; Anchorage stands out for Oil, Chemicals and Nonmetallic Minerals; and Seattle for Motor Vehicles (which includes aircrafts). It is also interesting to note that the two largest cities in the country are also top cities in several sectors, with New York dominating in most service sectors while Los Angeles stands out in several manufacturing sectors.

Table 4 compares estimates of productivity elasticities with respect to city size and employment composition obtained from our model inversion to those found in previous work. In particular, we report elasticities with respect to city size from the meta analysis carried out in Melo et al. (2009). As reported in Table 4, all results point to a positive relationship between TFP and city size. Moreover, our findings fall within the range of reduced form estimates found in the literature, with the possible exception of services. However, as in previous literature, the elasticity of (tradable) services productivity with respect to city size is substantially larger than that of manufacturing. To the extent that regional prices are not readily available, the relationship between TFP and city size estimated in some of the existing literature captures variations in nominal TFP, that is $\ln TFP_n^j + \ln P_n^j$. In other words, while our model inversion produces measures of P_n^j, the absence of local price data can otherwise bias downward empirically estimated elasticities of TFP with respect to city size. Indeed, our findings indicate that elasticities of real TFP with respect to city-size are somewhat larger than those of nominal TFP.

We also compare our TFP regressions coefficients with respect to the share of CNR workers to those estimated by Moretti (2004b) using a panel of firms (we use the CNR share of employment, whereas he uses the college educated share of employment). Again, our TFP measures are consistent with semi-elasticities that are of the same sign and comparable in magnitude to those in Moretti (2004b). As before, the regression coefficients become larger when deflated by the model-consistent regional price index.

Table 3: City with Top TFP for each Industry

Industry	MSA
Food and Beverage	San Francisco, CA
Textiles	Los Angeles, CA
Wood, Paper, and Printing	Minneapolis, MN
Oil, Chemicals, and Nonmetallic Minerals	Anchorage, AK
Metals	Los Angeles, CA
Machinery	Houston, TX
Computer and Electronic	San Jose, CA
Electrical Equipment	Los Angeles, CA
Motor Vehicles (Air, Cars, and Rail)	Seattle, WA
Furniture and Fixtures	Los Angeles, CA
Miscellaneous Manufacturing	Los Angeles, CA
Wholesale Trade	New York, NY
Transportation and Storage	New York, NY
Professional and Business Services	San Jose, CA
Other	Los Angeles, CA
Communication	New York, NY
Finance and Insurance	New York, NY
Education	New York, NY
Health	New York, NY
Accommodation	San Francisco, CA

City with top TFP_n^j in equation (9) by industry defined

Table 5 shows the coefficients in Table 4 for tradable sectors disaggregated by industry. We find that the positive relationship between TFP and city size holds uniformly across all tradable sectors. In addition, we find that the semi-elasticity of TFP in Computer and Electronics with respect to employment composition across cities is more than twice as large as the average for manufacturing, replicating the finding by Moretti (2004b) for high-tech sectors.

4.2 Estimating Production Externalities by Worker Type

So far, we have described the equilibrium levels of ocupation-specific productivity consistent with observed data on wages and sectoral employment, $T_n^{kj} \equiv \left(\lambda_n^{kj}\right)^{\gamma^j}$. Having obtained these productivity measures through the model inversion, we now turn to estimating their relationship to the scale and population composition of cities.[20]

We assume that occupational spillovers have the same labor augmenting effect across sectors. This assumption is consistent with Ellison et al. (2010) who find that aside from natural advantages and input-output linkages, occupational complementarities are the main source of industrial colocation. Hence, we let

$$\ln T_n^{kj} = \tau^{R,k}\gamma^j \ln\left(\frac{L_n^k}{L_n}\right) + \tau^{L,k}\gamma^j \ln\left(L_n\right) + \ln \widehat{T}_n^{kj}, \tag{10}$$

where $\widehat{T}_n^{kj}$ is an exogenously determined component of technology. In turn, this term is given by

$$\ln \widehat{T}_n^{kj} = a_0^k + a_Z^k Z_n^j + d^{kj} + u_n^{kj},$$

[20]The empirical exercise focuses on tradable sectors, for which our model generates measures of productivity separate from local housing supply.

Table 4: Elasticities of TFP with respect to City Size and Employment Composition

	$\ln(L_n)$		$\frac{L_n^{CNR}}{L_n}$	
	Real	Nominal	Real	Nominal
Average[1]	0.039	0.027	0.807	0.630
Manufacturing Average	0.028	0.017	0.581	0.458
Tradable Services Average	0.053	0.039	1.083	0.840
Melo et. al. Economy[2]	0.031 (0.099)		—	
Melo et. al. Manufacturing	0.040 (0.095)		—	
Melo et. al. Services	0.148 (0.148)		—	
Moretti College Share[3] (Manufacturing)	—		0.846 (0.102)	

Average coefficients from univariate OLS regression of $\ln TFP_n^j$ defined in equation (9) on $\ln L_n$ and L_n^{CNR}/L_n.

1. Excludes non-tradables

2. Melo et al. (2009) "A Meta-analysis of estimates of urban agglomeration economies", Table 2. "By type of response variable" and "By industry group."

3. Moretti (2004b) "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions", Table 2. College share in other industries, Cobb-Douglas production, 1992.

where Z_n^j is a vector of observable city/industry characteristics, d^{kj} denotes a set of industry dummies, and u_n^{kj} captures unobserved city-specific sources of natural advantages in the production of sector j goods with workers of type k.

Equation (10) allows for an agglomeration effect that depends on city size, through $\tau^{L,k}$, and an additional effect related to the share of each worker type, through $\tau^{R,k}$. The elasticity of productivity with respect to the agglomeration of a given type k is[21]

$$\frac{\partial \ln \lambda_n^{kj}}{\partial \ln L_n^k} = \tau^{R,k}\left(1 - \frac{L_n^k}{L_n}\right) + \tau^{L,k}\frac{L_n^k}{L_n}. \tag{11}$$

Therefore, when $\tau^{L,k} < \tau^{R,k}$, individuals in a given occupation, k, have a larger marginal effect on that occupation in cities where those individuals are less represented. However, there are also cross-occupational effects. Specifically, for $k \neq k'$, we have that

$$\frac{\partial \ln \lambda_n^{kj}}{\partial \ln L_n^{k'}} = -\left(\tau^{R,k} - \tau^{L,k}\right)\frac{L_n^{k'}}{L_n}, \tag{12}$$

which implies negative cross-occupational externalities when $\tau^{L,k} < \tau^{R,k}$. This congestion effect increases with the share of workers in alternative occupations.

The first column of Table 6 reports the coefficients from a simple OLS regression where we allow for two-way clustered standard errors by city and industry. These coefficients are positive and significant. They indicate that individual productivity is enhanced by the presence of other workers

[21] See Glaeser and Gottlieb (2008) for a discussion of the marginal implications of those elasticities. Note also that the specification can be readily extended to more occupations as long as we restrict all cross-occupation effects to take place, directly or indirectly, through total population.

Table 5: Sectoral Elasticities of TFP with respect to city size and employment composition

	$\ln(L_n)$		$\frac{L_n^{CNR}}{L_n}$	
	Real	Nominal	Real	Nominal
Food and Beverage	0.024	0.014	0.433	0.315
Textiles	0.029	0.007	0.293	0.251
Wood, Paper, and Printing	0.024	0.012	0.575	0.292
Oil, Chemicals, and Nonmetalic Minerals	0.047	0.031	0.933	0.797
Metals	0.024	0.010	0.488	0.265
Machinery	0.018	0.014	0.428	0.368
Computer and Electronic	0.056	0.043	1.480	1.219
Electrical Equipment	0.016	0.013	0.406	0.358
Motor Vehicles (Air, Cars, and Rail)	0.023	0.014	0.539	0.389
Furniture and Fixtures	0.018	0.009	0.177	0.284
Miscellaneous Manufacturing	0.033	0.020	0.640	0.503
Wholesale Trade	0.045	0.037	0.916	0.822
Transportation and Storage	0.032	0.025	0.602	0.501
Professional and Business Services	0.054	0.036	1.145	0.829
Other	0.057	0.046	1.073	0.945
Communication	0.045	0.042	0.995	0.939
Finance and Insurance	0.062	0.051	1.358	1.140
Education	0.076	0.052	1.632	1.104
Health	0.065	0.036	1.424	0.787
Accommodation	0.039	0.027	0.598	0.494

Coefficients from univariate OLS regression of $\ln TFP_n^j$ defined in equation (9) on $\ln L_n$ and L_n^{CNR}/L_n.

of the same occupational group. The coefficients also indicate the presence of congestion effects since cross-occupational externalities are negative. These OLS estimates, however, are potentially biased since workers of a given type may choose to live in cities where they are relatively most productive. This would induce a correlation between the exogenous component of worker productivity, $\widehat{T}_n^{kj}$, and the share of each type of worker in a given city. Moreover, the estimates might be biased if there are omitted variables which are correlated with both $\widehat{T}_n^{kj}$ and the occupational ratio.

To help address the omitted variable bias, Table 6 explores the effects of adding various controls to our basic OLS regression. Column 2 includes dummies for 9 census divisions interacted with industry dummies.[22] These should absorb many of the geographical and historical components that may jointly determine amenities and productivity in different places. Column 3 introduces geographic amenities constructed by the United States Department of Agriculture (USDA) that include measures of climate, topography and water area.[23] These controls allow for the possibility that the same geographic characteristics that may lead workers to choose certain cities may also influence their productivity. Column 4 introduces the share of manufacturing workers in 1920 as a control.[24] This aims to extract long standing factors that may influence the industrial composition in individual places. Finally, column 5 adds controls for demographic characteristics of different cities, including racial composition, gender split, the fraction of immigrant population,

[22] They are 1. New England, 2. Mid-Atlantic, 3. East North Central, 4. West North Central, 5. South Atlantic, 6. East South Central, 8. Mountain and 9. Pacific

[23] Geographic controls include average temperature for January and July, hours of sunlight in January, humidity in July from 1941 to 1970, variation in topography, and percent of water area.

[24] Just as with our labor force variables of interest, this and other controls are likewise interacted with the value added shares γ_n^j.

and age structure.[25] Together, these controls help narrow down the identification of the externality coefficients to the extent that more productive cities attract individuals of certain demographic make-up.

The point estimates of the coefficients on CNR workers change only slightly with the controls, while they increase the effect of labor market composition on non-CNR workers. These controls help extract exogenous sources of productivity variation that affect individual location decision. However, any residual variation in productivity may still be correlated with population levels and composition. In order to further account for those residual effects, we adopt an instrumental variable strategy, drawing on the existing empirical literature for candidate sources of exogenous variation. The key difference here is that we seek to explain productivity measures extracted from a structural model directly.

[25]Demographic controls are, by city, the percent female, black, hispanic, and percent in the age bins 16-25 and 26-65 (observations related to the younger than 16 population are dropped from the sample, and the age bin 66+ is omitted from the regression).

Table 6: OLS Estimates

	(1)		(2)		(3)		(4)		(5)	
VARIABLES	CNR	non-CNR	CNR	non-CNR	CNR	non-CNR	CNR	non-CNR	CNR	non-CNR
$\gamma_n^j log(\frac{L_n^k}{L_n})$	0.819***	0.631***	0.810***	0.672***	0.836***	0.691***	0.834***	0.670***	0.889***	0.702***
	(0.13)	(0.23)	(0.12)	(0.20)	(0.12)	(0.22)	(0.12)	(0.21)	(0.12)	(0.22)
$\gamma_n^j log(L_n)$	0.422***	0.345***	0.419***	0.344***	0.409***	0.346***	0.410***	0.343***	0.386***	0.322***
	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)	(0.04)
Jan. Temp					0.0213	-0.0556**	0.0130	-0.0590**	0.00557	-0.0439**
					(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.02)
Jan. Hrs Sun					0.0747***	0.0466***	0.0730***	0.0477***	0.0671***	0.0570***
					(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
July Temp					-0.0813***	-0.0266	-0.0750***	-0.0208	-0.0693***	-0.0397**
					(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
July Humid					-0.0213	-0.0157	-0.0150	-0.0125	0.000915	-0.0261
					(0.02)	(0.03)	(0.02)	(0.03)	(0.02)	(0.02)
Topography					-0.0270*	-0.0246**	-0.0258*	-0.0239**	-0.0349***	-0.0177
					(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
% Water Area					0.0364***	0.00390	0.0371***	0.00426	0.0342***	0.00868
					(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
% 1920 Mfg Workers							-0.00420	0.00787	-0.000532	0.00434
							(0.01)	(0.01)	(0.01)	(0.01)
% female									-0.0547***	-0.0493***
									(0.01)	(0.01)
% black									0.00217	0.0379***
									(0.02)	(0.01)
% hispanic									0.0265*	0.00226
									(0.02)	(0.01)
% Age 16-25									0.00983	0.0407**
									(0.01)	(0.02)
% Age 26-65									0.0346**	0.0662***
									(0.01)	(0.02)
Industry FE	X	X	X	X	X	X	X	X	X	X
Census Division FE			X	X	X	X	X	X	X	X
Observations	7,640	7,640	7,640	7,640	7,560	7,560	7,460	7,460	7,460	7,460
R-squared	0.607	0.770	0.658	0.804	0.668	0.811	0.669	0.814	0.672	0.817

Regressions estimates equation (10). Dependent variable is $\ln T_n^{kj}$ obtained from model inversion procedure described in text. Standard errors in parentheses, clustered two-ways by city and by industry. *** p<0.01, ** p<0.05, * p<0.1

4.2.1 Instrumenting for Employment Levels and Composition

In order to isolate the residual simultaneity between exogenous productivity variation and labor allocation, we resort to variants of instruments proposed in the literature. Specifically, we follow Ciccone and Hall (1996) and use population a century ago to capture historical determinants of current population. We also follow Card (2001) and Moretti (2004a), and use variation in early immigrant population and the presence of land-grant colleges to capture historical determinants of skill composition across cities. A detailed discussion of the particular instruments is provided in Appendix B.

Table 7 shows the estimation results with instrumental variables and all the controls. The first column repeats the OLS results in the last column of Table 6, and the second column shows the corresponding two-stage-least-squares estimates. Those are similar to the OLS estimates, being well within one standard error from one another. To evaluate the strength of the instrumental variables, we follow the procedure in Sanderson and Windmeijer (2016) to obtain separate first-stage F statistics for each of the endogenous variables.[26] Since the F-statistics are below the value of 10 recommended by Staiger et al. (1997), the estimates may have some bias and incorrect standard errors. The literature on IVs then recommends the use of limited information maximum likelihood (LIML) estimators. The third column of Table 7 carries out the estimation using a continuously updated GMM estimator (GMM-CUE), similar to a limited information maximum likelihood estimator but which allows for clustered and heteroskedastic standard errors. The Stock and Yogo (2005) critical values in the LIML model for a p-value of 5% to be 10% or better is 6.46, at or close to our obtained values.

These instruments, inspired from the previous work mentioned above, exploit the idea that after allowing for controls, all long-term effects of either historical immigration enclaves, land-grant college location, or historical population on local productivity derive from their impact on current occupational composition and population. To verify that our empirical strategy indeed identifies external effects, we carry out the same regressions on data generated by a counterfactual allocation in which we set $\tau^{R,k} = \tau^{L,k} = 0$ for all occupations k. The results are presented in Table 10 in the Appendix. They confirm that the OLS estimates for the effect of occupation shares on productivity are biased downward (the coefficients are now negative), whereas the estimates for the external productivity effects related to population exhibit little bias. More importantly, the exercise also shows that our IV's successfully eliminate most of those biases, especially so in the GMM-CUE estimates.

As a final measure of robustness, we carry out an estimation exercise using IV's implied by the model. Recall that, in our framework, the size and composition of population in different locations is determined simultaneously by local productivity, amenities, input-output linkages, and trade costs. Thus, we construct a counterfactual allocation where, for each industry and occupation, we set productivity to be fully exogenous and equal to the averages, across cities, of the productivity parameters, T_n^{kj}. We then use the counterfactual employment shares and totals implied by that exercise as instruments. The results, presented in Table 11 in the Appendix, confirm the main findings. Namely, the effect of population composition is larger than that of city size, and the compositional effect is larger for CNR workers than for non-CNRs.

[26] This follows largely the intuition laid out by Angrist and Pischke (2008), that requires strong IVs to predict the two endogenous variables independently from one another.

Table 7: Instrumental Variables Estimate

	(1) OLS		(2) 2SLS		(3) CUE	
VARIABLES	CNR	non-CNR	CNR	non-CNR	CNR	non-CNR
$\gamma_n^j log(\frac{L_n^k}{L_n})$	0.889***	0.702***	1.177***	0.263	1.304***	0.835*
	(0.12)	(0.22)	(0.38)	(0.51)	(0.38)	(0.51)
$\gamma_n^j log(L_n)$	0.386***	0.322***	0.334***	0.284***	0.349***	0.357***
	(0.05)	(0.04)	(0.06)	(0.05)	(0.06)	(0.04)
Observations	7,460	7,460	7,460	7,460	7,460	7,460
K.P. F			3.912	5.425	3.912	5.425
S.W.F. L_n^k Share			5.975	8.369	5.975	8.369
S.W.F. L_n			5.997	8.587	5.997	8.587

Regressions estimates equation (10). Dependent variable is $\ln T_n^{kj}$ obtained from model inversion procedure described in text. Standard errors in parentheses, clustered two-ways by city and by industry. *** p<0.01, ** p<0.05, * p<0.1

4.3 The Role of Externalities in Spatial Occupational Polarization

We adopt the GMM-CUE coefficients in the last column of Table 7 as our benchmark. These coefficients imply that cross-occupational externalities are negative for CNR workers in that $\tau^{L,CNR}$ is significantly smaller than $\tau^{R,CNR}$. Hence, non-CNR workers create negative congestion effects for CNR workers. In contrast, for non-CNR workers, the difference between $\tau^{R,nCNR}$ and $\tau^{L,nCNR}$ is not significant indicating no clear evidence of congestion effects from CNR workers to non-CNR workers.[27]

The externality effects coming from the local occupational composition are also clearly substantial. They imply, all else equal, that moving from Winston-Salem, NC with a share of CNR employment of 36 percent, corresponding to approximately the 76th percentile of the distribution of CNR shares, to Austin, TX, with share of 42 percent, closer to the 95th percentile, increases T_n^{kj} for CNR workers by approximately 6 percent and reduces that of non-CNR workers by close to 3 percent. Agglomeration externalities are similarly important. Moving from a city near the 75th percentile, such as Trenton, NJ, with approximately 182 thousand workers, to one near the 88th percentile such as Rochester, NY, with approximately 433 thousand, would imply a gain of close to 11% for both types of workers.

In Section 2 we conjectured that occupational externalities account for salient patterns in the data related to the polarization of occupations and wage inequality across cities. Given the model, the recovered TFP measure obtained through its inversion, as well as the externalities in TFP we estimated, we now verify this basic intuition.

Counter-factual equilibrium allocations generated in the absence of externalities (i.e., $\tau^{R,k} = \tau^{L,k} = 0$ for all k) are presented in the blue dots in Figures 5 and 6. Figure 5 shows that absent externalities, the relationship between the share of CNR workers and the wage premium indeed becomes negative, indicating that the relative abundance of CNR workers now decreases their relative compensation. Furthermore, Figure 6 shows that without externalities, the equilibrium wage bill share increases less than one-for-one, as opposed to more than one-for-one, with the wage bill share predicted by the industrial composition across cities. To see why, note that absent production externalities induced by employment size and composition, productivity is pinned down

[27] The p-value when testing the hypothesis that $\tau^{R,k} - \tau^{L,k}$ is positive is equal to 0.017 for the case of CNRs and equal to 0.34 for the case of non-CNRs.

exogenously. Therefore, CNR workers in a city that has a comparative advantage in the production of CNR intensive goods will generally earn higher wages. Firms in that city will consequently substitute CNR workers for non-CNR workers and, given an elasticity of substitution between occupations greater than 1, see a reduction in its CNR wage share.

These exercises point to the patterns identified in Section 2 as being effectively driven by occupation-specific elasticities. Given the significance of these externalities, the optimal and equilibrium allocations differ. This in turn creates a role for optimal spatial policy to which we turn next.

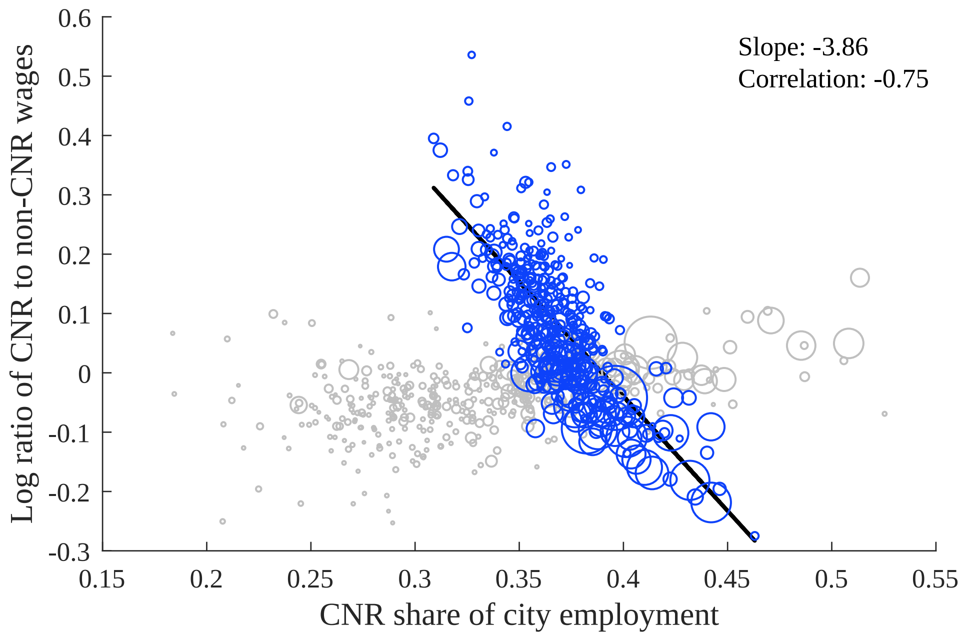

Figure 5: Occupational share and wage premium - no externalities

Counterfactual values obtained from assuming no externalities ($\tau^{R,k} = \tau^{L,k} = 0$), while keeping the exogenous part of productivity as originally quantified (blue dots). Grey dots correspond to equilibrium values in Figure 1

.

5 Optimal Allocation

We now describe the optimal allocation and the policies that implement it. We start by defining social preferences and setting up the planner's problem.

5.1 The Planner's Problem

The planner's problem takes as given that workers in each occupation can freely move across cities. Under this assumption, the expected utility of a worker of type k is given by

$$v^k = \Gamma\left(\frac{\nu-1}{\nu}\right)\left(\sum_n \left(A_n^k C_n^k\right)^{\nu}\right)^{\frac{1}{\nu}}.$$

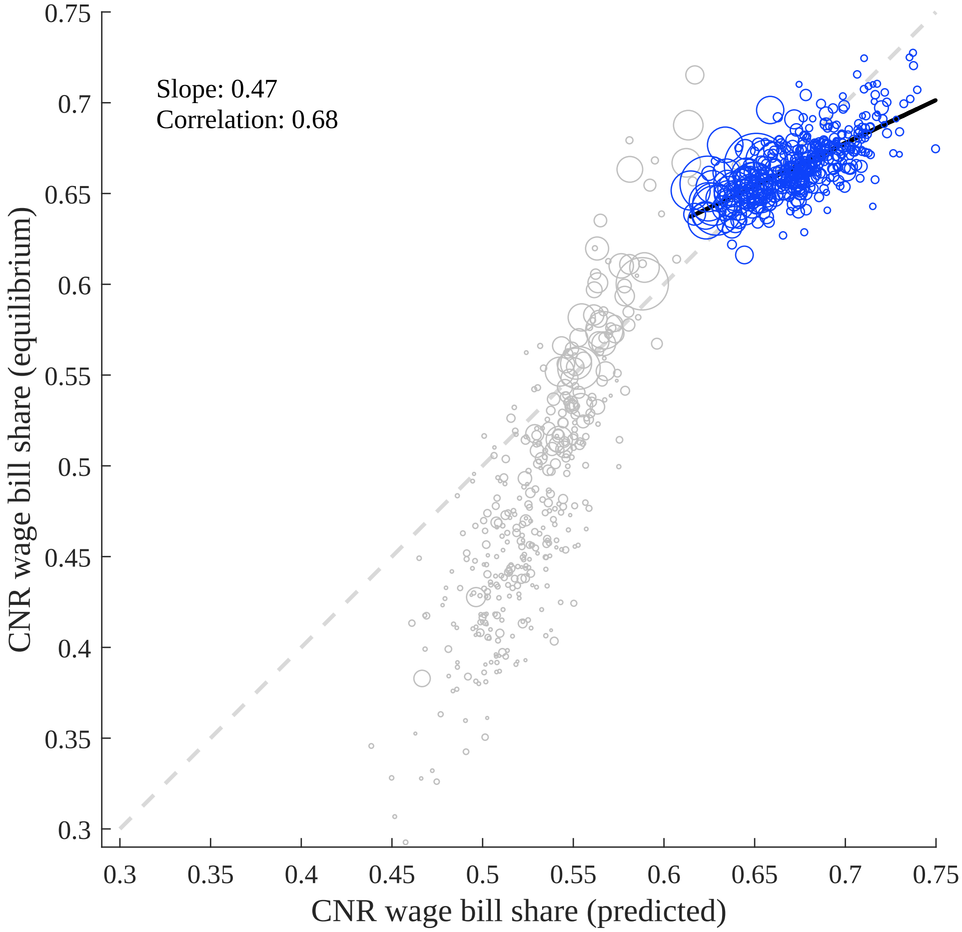

Figure 6: Occupational wage bill share: predicted vs equilibrium - no externalities

Counterfactual values obtained from assuming no externalities ($\tau^{R,k} = \tau^{L,k} = 0$), while keeping the exogenous part of productivity as originally quantified (blue dots). Grey dots correspond to equilibrium values depicted in Figure 3.

Then, if ϕ^k denotes the welfare weights for each occupation, we can postulate the generalized social welfare function

$$\mathcal{W} = \sum_k \phi^k U \left[\Gamma\left(\frac{\nu-1}{\nu}\right) \left(\sum_{n=1}^{N} \left(A_n^k C_n^k\right)^{\nu} \right)^{\frac{1}{\nu}} \right] L^k. \tag{13}$$

where U is an increasing and concave function.[28] The planner maximizes the expression in (13) subject to the availability of labor in each city and occupation (2), the constraints on the use of labor in each occupation and city (5), the constraints on the use of land and structures (6), the resource constraints associated with final goods in each city and sector (7), the resource constraints associated with intermediate goods across all varieties $\mathbf{z}$ in each industry j and city n (8), and the constraints that household consumption of different goods and input flows be non-negative.

[28]This generalized social welfare function nests the leading cases of a utilitarian planner, in which case U is linear, and the limit in which U becomes infinitely concave and so $\mathcal{W}$ approximates the max-min welfare function of a Rawlsian planner.

The key difference between the optimal and equilibrium allocations stems from a wedge between the private and social marginal products of labor. Lemma 1 characterizes this wedge.

Lemma 1. *Let Δ_n^k denote the wedge between the private and the social marginal value of a worker in occupation k in city n. Then*

$$\Delta_n^k = \sum_{k',j} w_n^{k'} \frac{L_n^{k'j}}{L_n^k} \frac{\partial \ln \lambda_n^{k'j}(\mathbf{L}_n)}{\partial \ln L_n^k}. \tag{14}$$

This expression for the wedge points to the distortions that the planner seeks to correct. In particular, it is increasing in the elasticity of worker productivity with respect to the number of workers in occupation k, $\partial \ln \lambda_n^{k'j}(\mathbf{L}_n)/\partial \ln L_n^k$, in all industries and occupations. Moreover, the contribution of this elasticity in a given industry-occupation, for a given city, varies with the proportion of workers in that industry and occupation, $L_n^{k'j}/L_n^k$, as well as their private marginal product or wage, w_n^k.

To focus ideas further, consider our case with two occupations, k and k', where spillover elasticities are the same in all sectors. In that case we have that

$$\Delta_n^k = w_n^k \frac{\partial \ln \lambda_n^k(\mathbf{L}_n)}{\partial \ln L_n^k} + w_n^{k'} \frac{L_n^{k'}}{L_n^k} \frac{\partial \ln \lambda_n^{k'}(\mathbf{L}_n)}{\partial \ln L_n^k}.$$

The estimated externality parameters in Table 7 imply, using equations (11) and (12), that occupation k's own elasticity, $\partial \ln \lambda_n^k(\mathbf{L}_n)/\partial \ln L_n^k$, is positive while the cross-elasticity, $\partial \ln \lambda_n^{k'}(\mathbf{L}_n)/\partial \ln L_n^k$, is negative. It follows that the wedge for any given occupation k increases with its wage and decreases with the wage of occupation k'. Hence, the planner would like to increase the concentration of workers of a given occupation in places where those workers are most productive, and in places where workers in other occupations are less productive. The latter effect is stronger for CNRs than non-CNRs since $\partial \ln \lambda_n^{k'}(\mathbf{L}_n)/\partial \ln L_n^k$ is substantially larger when $k = \text{CNR}$. The end result is an increase in spatial polarization.

Given these wedges, the optimal policy is then most intuitively framed in terms of a set of taxes and subsidies that incentivize workers to move to cities where their spillovers are larger. Put another way, the planner internalizes the wedge between the private and social marginal productivity of workers. At the same time, a utilitarian planner also attempts to balance gains between different type of workers. Proposition 1 provides an exact characterization of this spatial policy.

Proposition 1. *If the planner's problem is globally concave, the optimal allocation can be achieved by a set of taxes and transfers such that*

$$P_n C_n^k = (1 - t_L^k)(w_n^k + \Delta_n^k) + \chi^k + R^k,$$

where $t_L^k = \frac{1}{1+\nu}$, and R^k is such that

$$\phi^k U'(v^k) v^k L^k = \sum_n P_n C_n^k L_n^k.$$

The proposition generalizes a key insight in Fajgelbaum and Gaubert (2018) to a multi-industry environment. Because spillover elasticities are not industry-specific, one need not keep track of sectoral employment shares in order to determine the optimal subsidy. Therefore, even with multiple industries, the optimal policy collapses to the special case in which occupational shares and wages become sufficient statistics.[29] Unlike Fajgelbaum and Gaubert (2018), however, it remains that the optimal allocation does depend on local industry-specific productivity (since they determine local employment and wages in both occupations), and has implications for the composition of industries across space. We show below that under the optimal policy, large and small cities expand industries in which a large share of their employment already resides, intensive in CNR and non-CNR workers respectively, while medium-size cities generally diversify across industries.

The condition of global concavity depends, in general, on the concavity of the U function in the planner's objective. In our numerical results, we assume that U is sufficiently concave to guarantee that Proposition 1 holds. Then, the result tells us that the planner's solution for household consumption differs from that implied by their budget constraint in two ways. First, the planner's solution depends on the social marginal product of labor, given by $w_n^k + \Delta_n^k$, rather than its private counterpart. Second, in the planner's solution, consumption increases less than one-for-one with the (social) marginal product of labor. This second element is optimal because, given heterogeneity in preferences for locations, households that choose to live in lower wage cities do so because their marginal utility of consumption is higher in those cities.[30]

5.2 The Value of Social Wedges Across Cities and Occupations

The wedge between the social and private marginal product of labor, Δ_n^k, may be calculated for each city and occupation using equation 14. Figures 7 and 8 show the deviations of those wedges from their (employment weighted) means for CNR and non-CNR workers, respectively. The average wedge for CNR workers is itself fairly large, at \$56,358 dollars per worker, or 79 percent of the mean CNR wage. The average wedge for non-CNR workers is more modest and negative, at -\$6,777. The wedge of non-CNR workers is negative because their presence in a given city is associated with a reduction in its share of CNR workers which then lowers the productivity of those workers. On average, non-CNR workers generate a net congestion effect. Together, these values imply an average gain of \$63,135 from switching a non-CNR for a CNR worker. This large value is the result of the relative scarcity of CNR workers, implying that using them productively makes a substantive difference. Furthermore, this large gain indicates that education and migration policies that create and attract CNR workers can potentially have high social value. Here, however, we take the supply of CNR and non-CNR workers as given.

We find a positive correlation between Δ_n^{CNR} the wedge between the social and private value of CNR workers, and city size (0.42). In contrast, the correlation between this wedge and the CNR share across cities is close to zero (0.038). These findings indicate that, given the concavity of external effects and the fact that there are diminishing returns to each factor, high CNR cities already exploit CNR externalities to a large degree in the decentralized equilibrium. However, the correlation with the CNR share becomes positive (0.356) when weighted by city size. Thus, there nevertheless remain gains to be exploited in larger, CNR intensive, cities. Externalities from CNR workers appear to be particularly large in New York, Houston, and cities in California and much less pronounced in Florida and, more broadly, in the South and Mid-West (except, modestly, in Chicago).

The overall patterns for $\Delta_n^{non-CNR}$, the wedge between the social and private value of non-CNR workers, are more pronounced in Figure 8. In particular, there exists a clear negative relationship between the wedge of non-CNR workers and both city size and the CNR share across cities. The

[29] As indicated in equation (14), when spillover elasticities are industry-specific, the optimal subsidy requires that the planner keep track of sectoral employment shares but not the details of intersectoral linkages.

[30] Fajgelbaum and Gaubert (2018) show that the heterogeneity in preferences induces the same optimal tax as an isoelastic negative spillover in amenities. See Davis and Gregory (2020) for a critique of this argument.

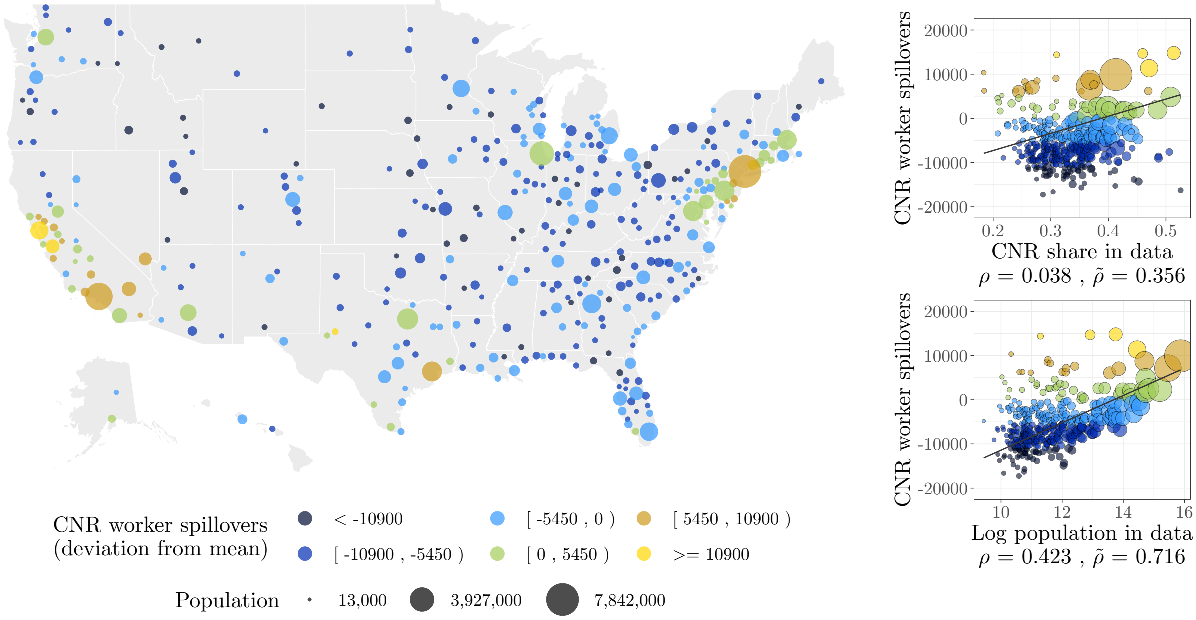

Figure 7: Δ_n^{CNR} (equilibrium values)

Δ_n^{CNR} captures the wedge between the social and the private marginal product of labor for CNR workers, as described in Lemma 1. Figure depicts deviation from employment weighted average (US$ 56,358). Each marker in the map refers to a CBSA. Marker sizes are proportional to total employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

social value of non-CNR workers relative to their private value is positive in many smaller cities. This wedge is also relatively large in some larger cities such as in Florida, Las Vegas, and Phoenix. These findings indicate that an optimal allocation would encourage or incentivize non-CNR workers to move to smaller non-CNR abundant cities. It is in those cities where they can make their largest contributions.

5.3 Quantifying the Optimal Allocation

In computing the optimal allocation, we set U to be concave enough such that the first order conditions are sufficient for optimality, and set the Pareto weights, ϕ^k, such that gains under the planner's solution are proportionately equal for both types of workers. Figures 9 and 10 show the percentage change in employment in the optimal allocation relative to the equilibrium allocation for CNR and non-CNR workers respectively. The results show that it is generally optimal for CNR workers to move to larger cities and for non-CNR workers to move to smaller cities, thereby exacerbating the spatial polarization of occupations. This increased spatial polarization follows from the spillover coefficient estimates in Section 4.2 which underscore that both types of workers (but particularly those in CNR occupations) become more productive when clustered with other workers of their own type.

As Figure 9 shows, increases in CNR workers under the optimal allocation are particularly large in cities like New York, San Francisco or San Jose, where the wedge between social and private marginal products of labor for CNR workers is especially large. These cities, together with other large cities including Chicago, Dallas, and Los Angeles, which are somewhat less CNR intensive, become cognitive hubs under the optimal allocation. More generally, the optimal policy creates

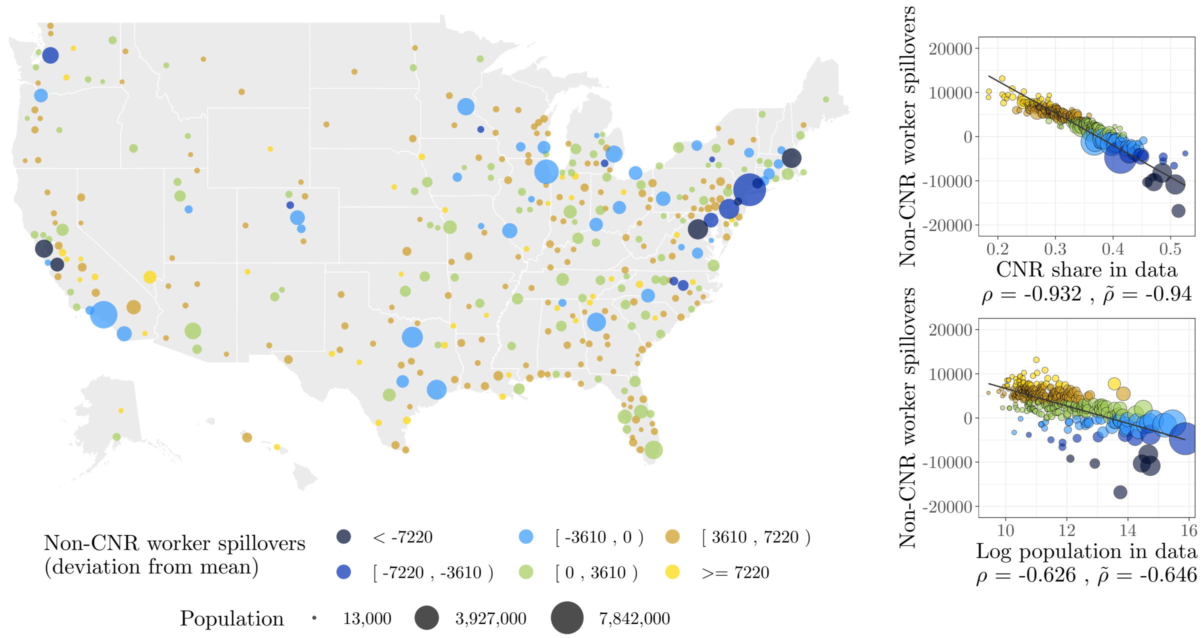

Figure 8: Δ_n^{nCNR} (equilibrium values)

Δ_n^{nCNR} captures the wedge between the social and the private marginal product of labor for non-CNR workers, as described in Lemma 1. Figure depicts deviation from employment weighted average (US$ -6,777). Each marker in the map refers to a CBSA. Marker sizes are proportional to total employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

cognitive hubs in larger cities that are already CNR abundant under the decentralized equilibrium. Given that trade is costly, cities that gain CNR workers are somewhat uniformly distributed in space according to overall economic activity. They constitute cognitive hubs in that they absorb CNR workers and are now surrounded by smaller cities with more non-CNR workers.

Figure 10 illustrates that while the planner generally chooses to incentivize non-CNR workers to move from large cities, a few large cities do nevertheless become more non-CNR abundant under the optimal allocation. This is the case for cities such as Miami, Las Vegas, Phoenix, and San Antonio where non-CNR workers have, in the decentralized equilibrium, a social marginal product that is larger than their private marginal product. These cities become new non-CNR centers. They specialize in non-CNR intensive industries, such as accommodation and retail, and grow in size since the inflow of non-CNR workers is larger than the exodus of CNR workers specified by the optimal allocation.

While the share of CNR workers increases in large cities under the optimal allocation, the top panel of Figure 11 also shows that these cities lose in overall population while smaller cities increase in size. The same pattern holds for cities with large and small CNR shares in the bottom panel of 11. New cognitive hubs, therefore, emerge along side growing small and non-CNR abundant cities. Put another way, the city size distribution evens out under the optimal allocation. This feature recognizes that while the productivity of CNR workers increases with the number of those workers, congestion also increases with city size. In particular, as discussed above, non-CNR workers generate negative congestion effects on the productivity of CNRs. Furthermore, heterogeneous location preferences imply that attracting the marginal CNR worker to a given city becomes increasingly difficult.

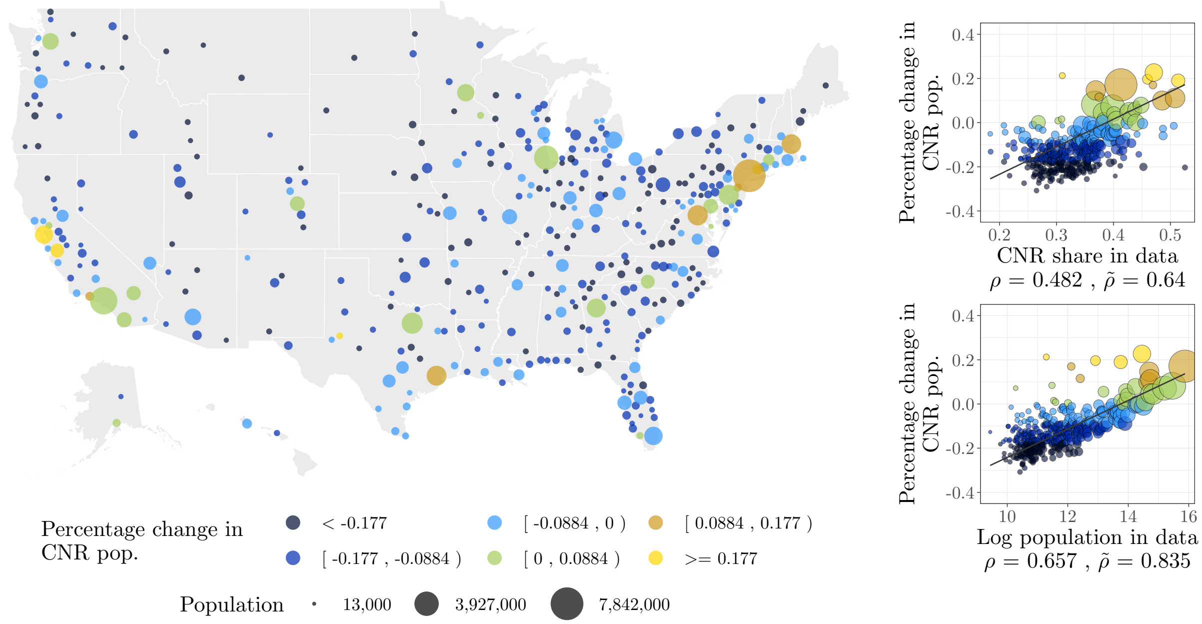

Figure 9: L_n^{CNR} (percentage change from data equilibrium)

Percentage change in employment of CNR workers between equilibrium and optimal values. Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

Along with cities becoming more even in size under the optimal allocation, we observe that both small and large cities generally increase their degree of industrial specialization, while medium-size cities tend to move towards greater industrial diversification. Figure 12 highlights this pattern using changes in the Gini coefficient associated with the distribution of wage bill shares across industries. The figure illustrates how these changes vary with the share of CNR employment and depicts a U-shaped pattern. In the efficient allocation, cities with low and high CNR worker shares become more specialized. In contrast, cities with intermediate CNR shares exhibit zero or negative changes in Gini coefficients, indicating no change or greater industrial diversification in those locations. This finding emerges because concentrating occupations is more valuable in cities that are particularly productive in industries intensive in a specific occupation. The planner's solution, therefore, prescribes further expanding industries intensive in either CNR or non-CNR occupations in cities that have a more extreme skill mix. Moreover, these cities tend to be at either end of the size distribution so that the U-shaped relationship shown in Figure 12 also holds, though somewhat attenuated, with respect to population (see Figure 24 in the Appendix).

As examples of how efficient allocations change the industrial composition landscape, Figure 13 highlights two cities at either end of the CNR share distribution. At one end, San Jose, CA, with close to 1 million workers, sees its share of CNR workers increase from 51.4% to 90.7%. This change reflects an increase in industrial specialization, summarized by a 0.11 change in the Gini coefficient and seen as an outward shift in its Lorenz curve in the left panel of Figure 13. It captures in part an increase of 19 percentage points in the employment share of San Jose's top industry, Professional and Business Services, and a 14 percentage point increase in that industry's wage bill share. At the other end, Harrisonburg, VA, with only 50,126 workers, sees instead its share of non-CNR workers increase from 73.0% to 83.0%. This change stems from the planner emphasizing employment in

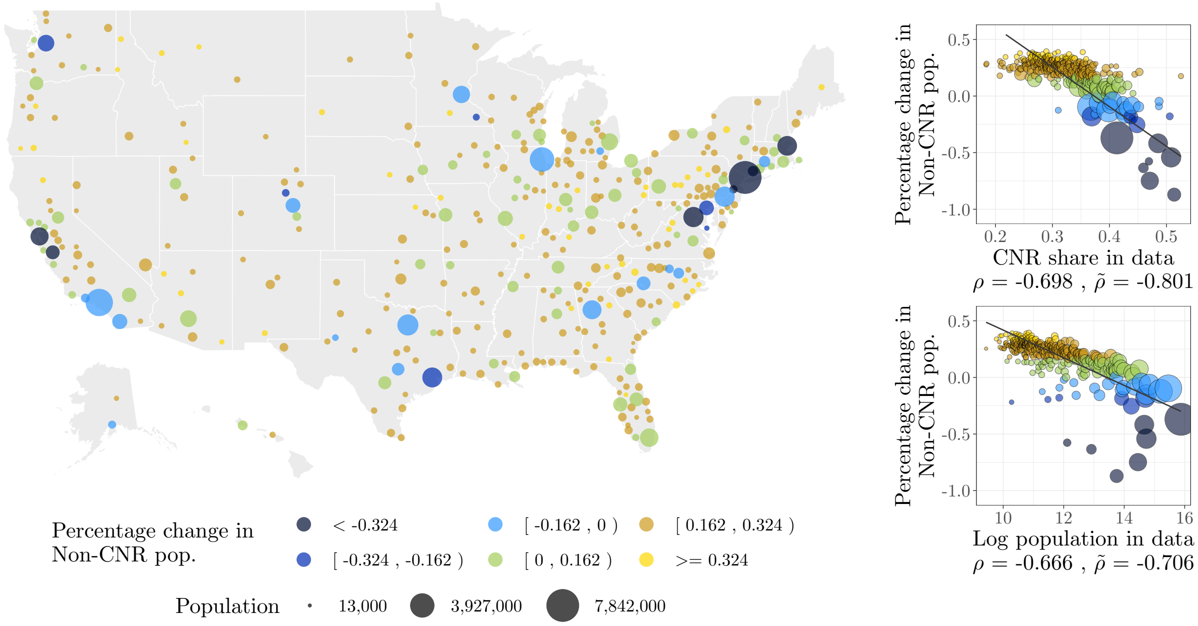

Figure 10: L_n^{nCNR} (percentage change from data equilibrium)

Percentage change in employment of non-CNR workers between equilibrium and optimal values. Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

the industries in which Harrisonburg's non-CNR workers are already intensively employed; hence, the Gini coefficient increases by 0.04. Under the optimal allocation, the employment share in Harrisonburg's top industry, Non-Tradables (Retail, Construction, and Utilities), increases by 2.0 percentage points while its wage bill sees a 3.0 percentage point rise.

In the middle of the city size distribution, changes in the Gini coefficient associated with wage bill shares across industries are generally close to zero or negative. In other words, many of those medium-size cities either stay with or diversify their industrial mix. The bottom panel of Figure 13, for example, shows an inward shift of the Lorenz curve for Scranton, PA. Scranton, a city of around 234,000 workers, sees its Gini coefficient fall by around 0.05 as its employment spreads out across more industries. The wage bill share of its top industry, namely Health, declines by 7.5 percentage points, while the wage bill share of its new largest industry, Non-Tradables, increases by 2.1 percentage points.

5.3.1 Taxes and Transfers: Implementing the Optimal Allocation

Implementing the optimal allocation involves transfers specific to each occupation and city. Including the transfers, total consumption is equal to total income in each city. These transfers serve several functions. First, they incentivize agents to move as described above. Namely, they incentivize CNR workers to move to cognitive hubs and non-CNR workers to move to smaller towns. Second, they guarantee that relative welfare gains are the same across occupations and locations. Thus, the planner compensates non-CNR workers for moving to smaller and less productive or amenable cities by implementing transfers from larger to smaller cities. These transfers in turn are mostly financed by CNR workers in larger cities. Note, however, that since CNR workers gain from the policy as well, they do not mind making the transfers. The net flow of resources received or

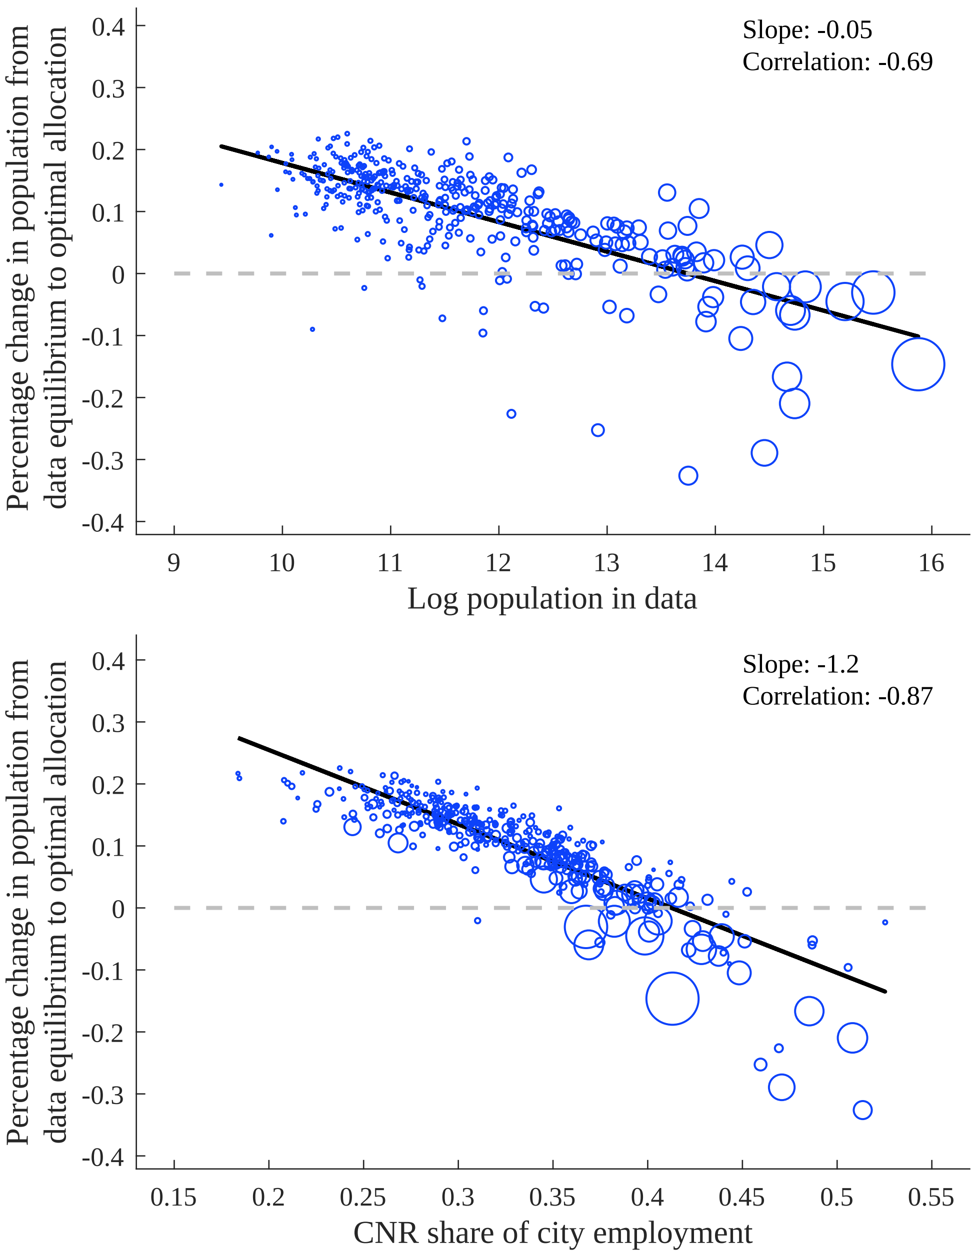

Figure 11: City size changes between data and optimal allocation

Each observation refers to a CBSA. Marker sizes are proportional to total employment.

paid by cities (i.e the trade balance) is shown in Figure 14, calculated as the difference between local nominal per capita consumption and output, $\left(\sum_k P_n C_n^k L_n^k - \sum_k w_n^k L_n^k - r_n H_n\right) / L_n$.

As expected, the trade balance of cognitive hubs is large and negative. Cities such as San Francisco and San Jose that are relatively large and very CNR intensive, specializing, respectively, in professional and business services and computer and electronics, send net payments of as much as \$40,834 per resident, while some of the smaller cities, such as Jacksonville, NC, specializing in accommodation and retail receive net transfers of close to \$18,000. These net transfers in some of the smaller cities amount to a form of basic income paid to all non-CNR workers in small cities

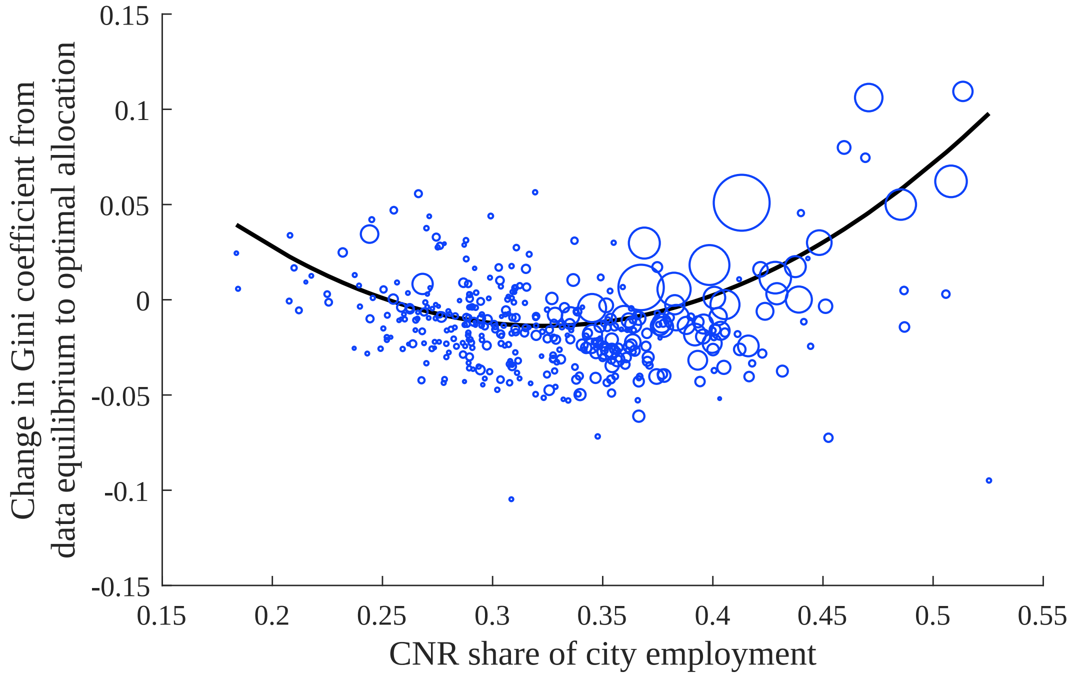

Figure 12: Changes in the Gini coefficient between the data and optimal allocation.

Each observation refers to a CBSA. Marker sizes are proportional to total employment. The solid-black line is a cubic fit on the data. The Gini is constructed using the Lorenz curves depicting within city wage bill and industry rank.

and financed (in net) by agents living in cognitive hubs. Again, there are some exceptions as some relatively large cities such as Las Vegas end up receiving net transfers since they become even larger centers of non-CNR employment.

The pattern of spillover coefficients we estimated, together with our model, implies that the optimal allocation exacerbates the extent of labor market polarization across space. However, this pattern does not reveal why labor markets are already as polarized as they are in the decentralized equilibrium. As indicated in Figure 4, larger cities are relatively more amenable to CNR workers, a pattern that survives even if we exclude the endogenous component of amenities highlighted by Diamond (2016). In Appendix E, we show that even with that component excluded, the planner chooses to create cognitive hubs by concentrating CNR workers in large cities.

The overall gains in welfare from implementing the optimal allocation amount to 0.59% of GDP for workers in both occupations. These gains are larger than those found in recent work by Bartelme et al. (2019) at 0.4% of GDP in a study of optimal industrial policy in a multi-country trade model. One reason why gains are not even larger in our setup is that the observed equilibrium allocation is already fairly polarized. In fact, in Section 6 we try to account for what has lead to this fairly polarized state starting from the prevailing conditions in 1980.

To achieve equal welfare gains across occupations, the optimal transfer scheme has two components. One that incentivizes agents to go to the 'right' location and is related to differences in Δ_n^k across locations. The other is a fixed transfer by occupation (R^k). This fixed transfer guarantees that all workers obtain equal gains from moving to the planner's allocation. This fixed transfer

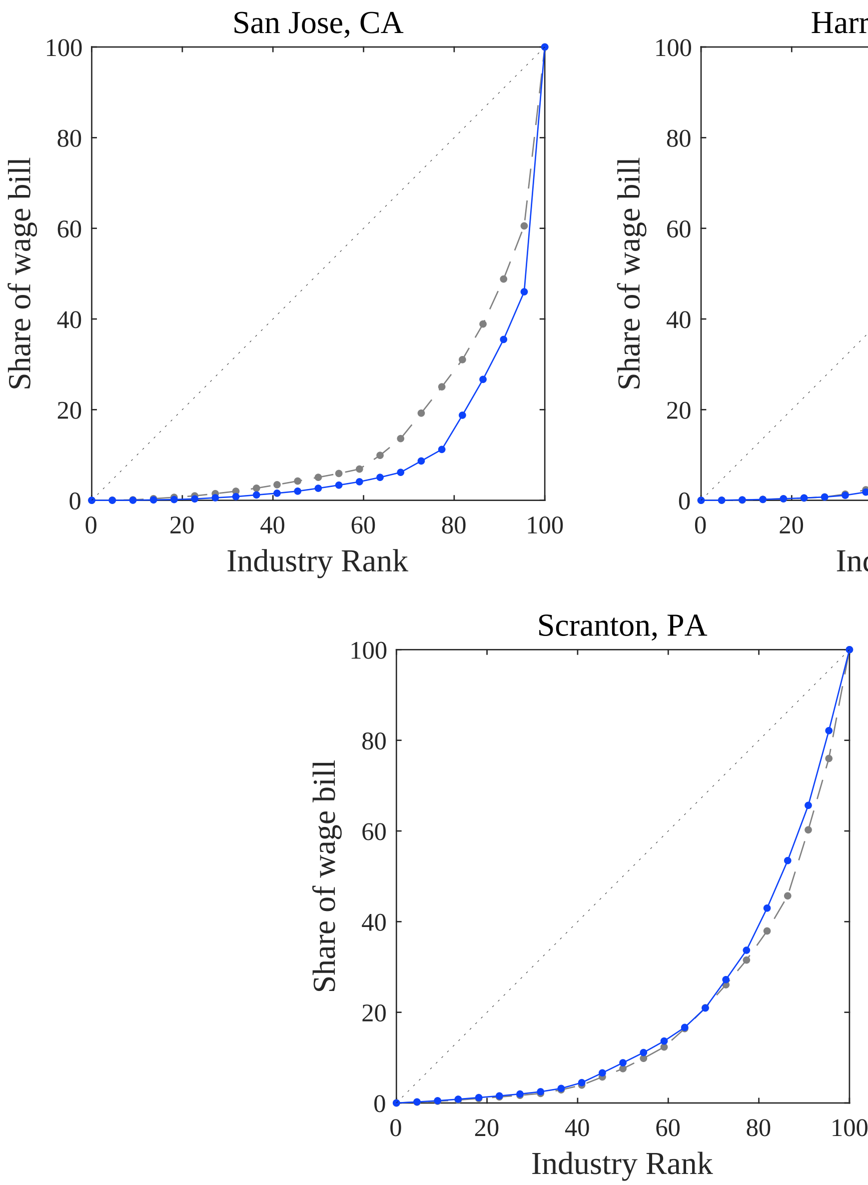

Figure 13: Shift in the Lorenz curve between the data and optimal allocation.

Each marker refers to an industry, where the dashed gray line shows the Lorenz curve for the data equilibrium and the blue line is for the optimal allocation.

amounts to a negative transfer (a payment) of -\$15,255 for CNR workers and a positive transfer of \$16,872 to all non-CNR workers. The latter can be implemented as a universal basic income that is paid by CNR workers. This transfer may then be considered as the redistribution that CNR workers are willing to accept to form the cognitive hubs where they can thrive.

The optimal transfers also involve a component that incentivizes CNR worker to move to large, CNR abundant cities and for non-CNR workers to move to smaller cities with smaller shares of CNR workers. This is achieved by giving large incentives to non-CNR workers to move out of large, and more markedly, CNR abundant cities. Due to externalities that are occupation-specific, this

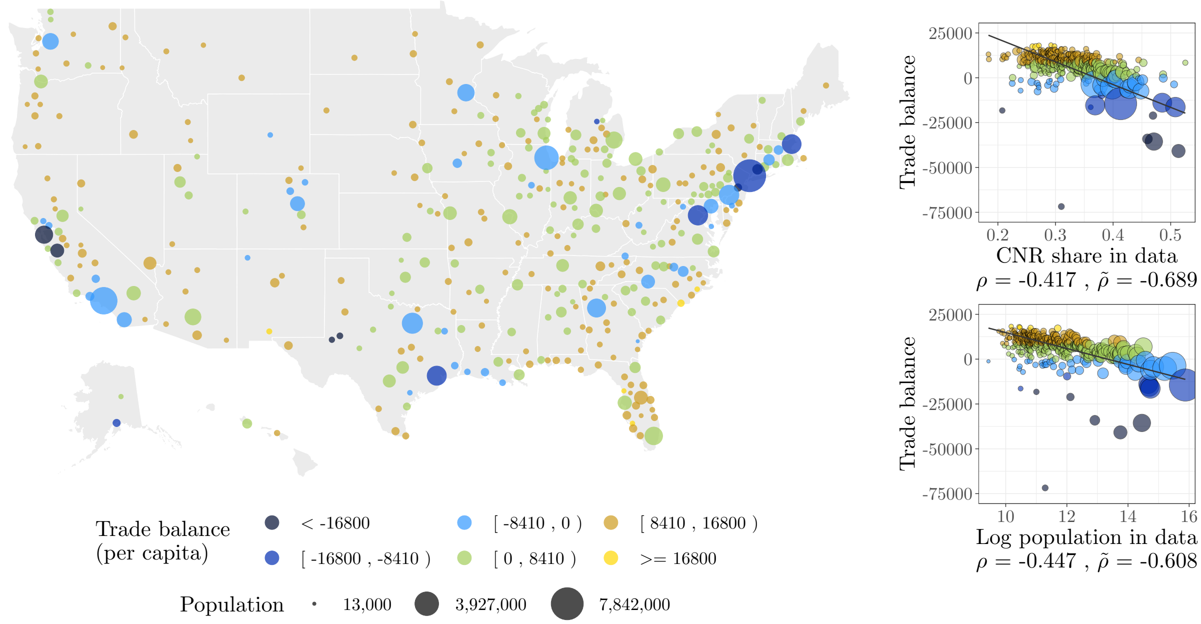

Figure 14: Trade balance per capita

Trade balance is defined as the difference in the optimal allocation between the value consumed and value added in each city ($\sum_k P_n C_n^k L_n^k - \sum_k w_n^k L_n^k - r_n H_n$). Trade balance per capita are given by those values divided by L_n. Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

reallocation yields larger CNR productivity increases in CNR abundant cities, which attracts CNR workers to these cities and eliminates the need for large net transfers to those workers. In fact, in a handful of the most CNR abundant cities, this effect is so strong that the planner prefers to balance it with negative transfers to avoid congestion.

The optimal transfers are, of course, related to the Δ_n^k wedges described in Figures 7 and 8, and further depicted in Figures 15 (for CNR workers) and 16 (for non-CNR workers).[31] In the median city, after incentives are taken into account, CNR workers contribute US$ 2,544. This number in part reflects the fact that CNR workers are socially valuable (recall that the wedge between the social and private value of CNR workers is as much as US$ 56,358). Note that CNR workers do not need to be particularly incentivized to stay in the large, CNR abundant, cities. In fact, as Figure 15 shows, transfers decrease slightly with CNR share and city size. The increases in productivity, and therefore wages, that result from the enhanced externalities in cognitive hubs is sufficient to attract these workers. CNR workers are net-recipients of transfers in only 8% of locations but net payers in 92% of locations. On the whole, payments from CNR workers range on net from US$ 309 (in the 10th percentile city) to US$ 4,422 (in the 90th percentile).[32]

Once incentive-based transfers are taken into account, non-CNR workers in the median city receive a net transfer of US$ 8,478, ranging from US$ 1,819 (in the 10th percentile city) to as much

[31] More specifically, total transfers are equal to $P_n C_n^k - w_n^k - \chi^k = \frac{\nu}{1+\nu}\Delta_n^k - \frac{1}{1+\nu} w_n^k + R^k$.

[32] Weighing by population net transfers to or from CNR workers range from a net contribution of close to US$ 113 (10th percentile) to a contribution close to US$ 10,000 (90th percentile), with the median net contribution being close to US$ 2,475.

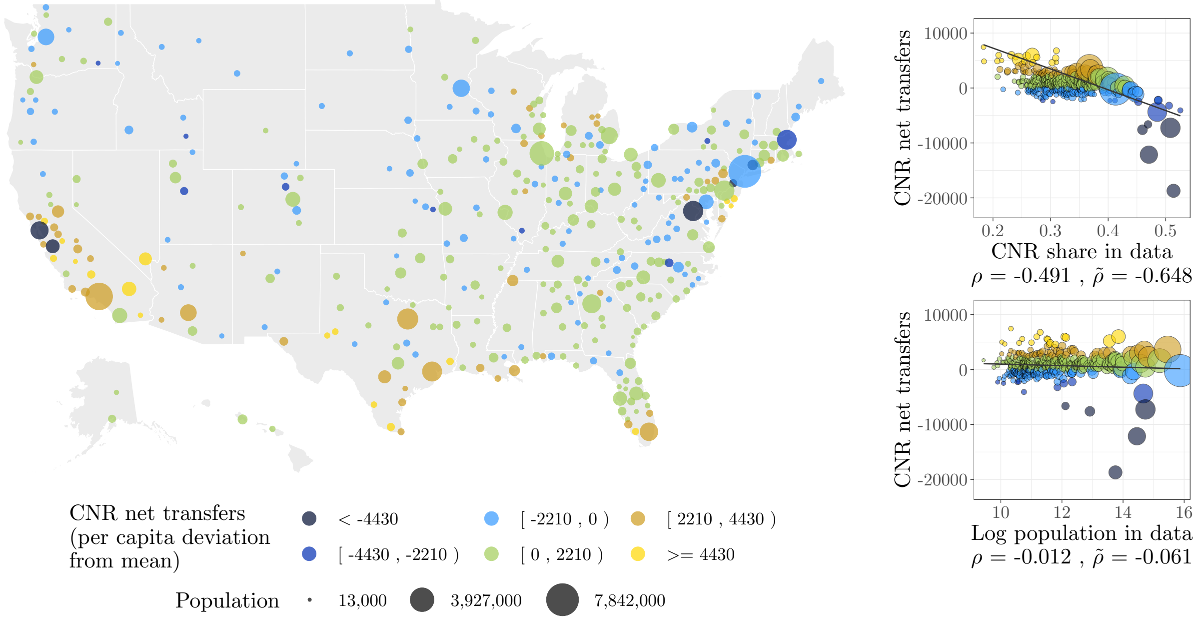

Figure 15: Optimal transfers to CNR workers (per CNR worker)

Optimal transfers per CNR workers are defined as the difference in the optimal allocation between the value consumed and the income they would receive given optimal wages and rents but absent the transfers ($P_n C_n^{CNR} - w_n^{CNR} - \chi^k$). Figure depicts deviation from employment weighted average (US$ -3,316). Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

as US$ 10,518 (in the 90th percentile city). Note that non-CNR workers have to be incentivized to move to smaller cities. However, because cognitive hubs offer a high wage to non-CNR workers (since those are the most productive cities while also ending up with fewer non-CNR workers), non-CNR workers in cognitive hubs pay a transfer as well to discourage other non-CNR workers from joining them. This accounts for the wide range in non-CNR transfers and reduces the average net burden on CNR workers.[33]

6 The Formation of Cognitive Hubs after 1980

The US economy has evolved towards the formation of cognitive hubs at least since the 1980s. Quantifying our model to 1980 data yields a set of fundamental characteristics of the economy that allows us to study this phenomenon in detail. In order to compare the spatial structure of the economy in 1980 to that in 2015, we want to abstract from aggregate technology trends. Thus, we first build a 'Baseline' economy that adds only *aggregate* changes in technology, population, and input shares to the 1980 economy.[34] The Baseline economy does not include any location-specific change in productivities across industries and occupations or in amenities. It also does not

[33] If we weight by population, transfers and contributions to non-CNR workers range from a net contribution of about US$ 6,599 (10th percentile) to net receipts of US$ 9,356 (90th percentile), with a median net receipt of about US$ 3,460.

[34] For details on the construction of this Baseline economy and other counterfactuals, see Appendix F.

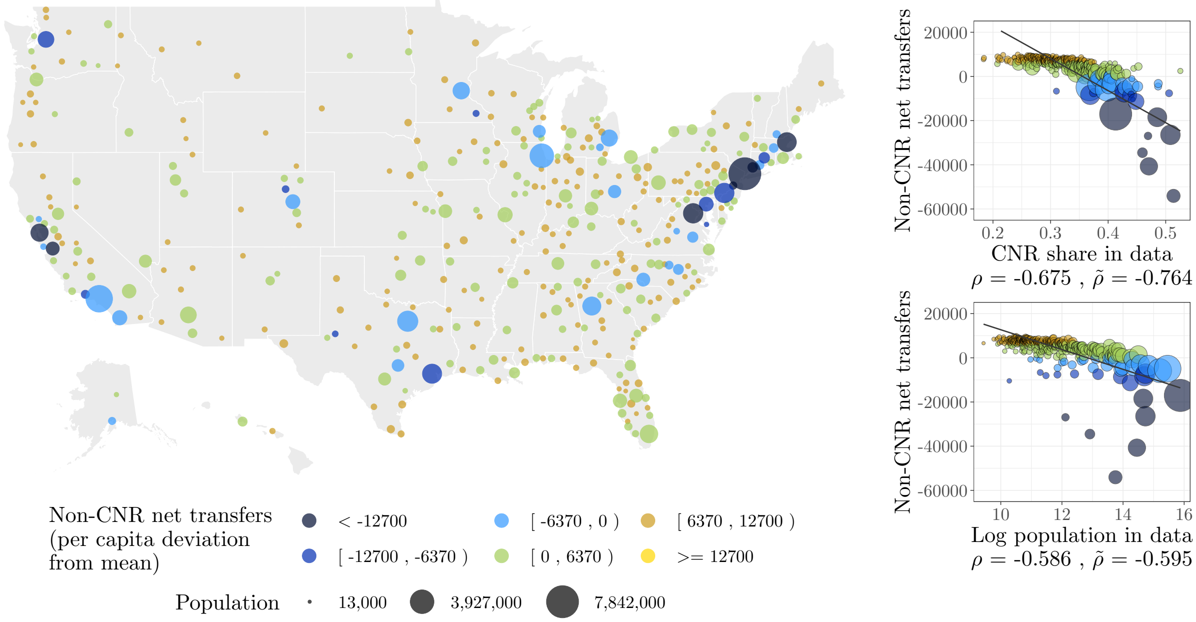

Figure 16: Optimal transfers to non-CNR workers (per non-CNR worker)

Optimal transfers per non-CNR workers are defined as the difference in the optimal allocation between the value consumed and value added in each city attributed to non-CNR workers ($P_n C_n^{nCNR} - w_n^{nCNR} - \chi^k$). Figure depicts deviation from employment weighted average (US\$ 2,020). Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

include changes in aggregate CNR shares of population. We study the role of these components by adding them gradually. As shown in Figure 17, in this Baseline 1980 economy, US population is concentrated in cities where the CNR share of employment was close to average. In the 2011-15 data, however, there is greater dispersion around the (now larger) average. Figure 17 further shows the distribution of CNR workers implied by the 2011-2015 planner's solution calculated above. Together, the histograms imply that the increasing concentration of CNR workers was in the direction implied by today's optimal allocation.

Given the move towards cognitive hubs over time, our model allows us to examine the forces underlying this evolution and obtain a quantitative assessment of their welfare relevance. We do this with a series of counter-factual exercises that clarify the importance of different forces in driving national trends. Thanks to the structural model, we can also do the same decomposition in the absence of externalities. Those exercises then provide us with a measure of the relevance of local spillovers for observed spatial trends. Table 8 shows how such a decomposition affects the welfare of CNR and non-CNR workers. The columns depict welfare levels relative to the Baseline 1980 economy for each occupation. The lower panel repeats the exercises for a world without externalities (i.e., where we set the externality elasticity parameters to zero).

The top row of the table shows that welfare in 1980 was lower for both groups and in all scenarios, as one would expect given underlying technological trends. The second row shows the effects of changing input shares. It is well documented that CNR intensive industries have become a larger part of the US economy, leading to relative gains for CNR workers as compared to non-CNR workers. The next step (rows 3 and 14) brings total population and average technology in

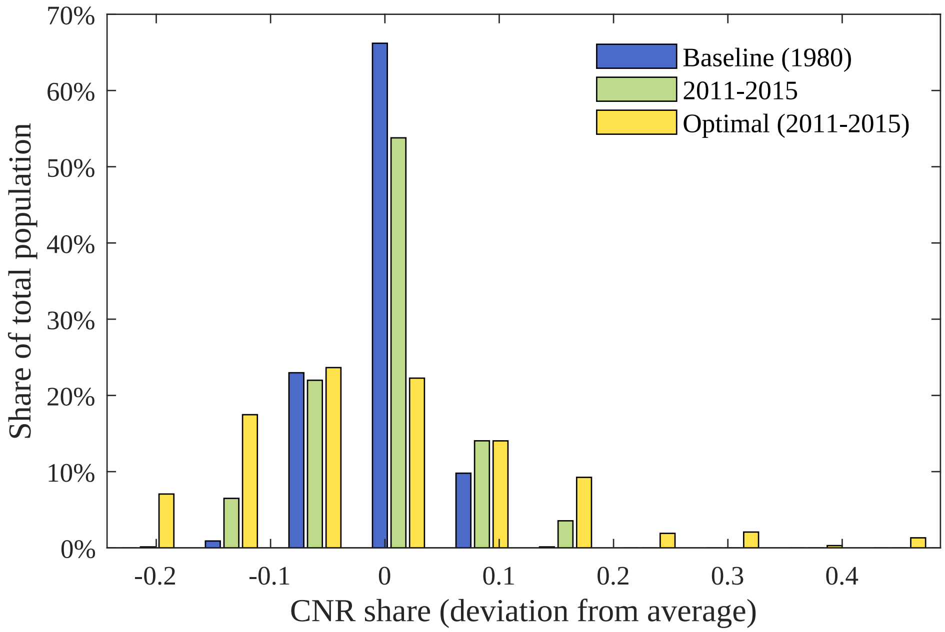

Figure 17: Distribution of population by CNR share in city of employment

Share of population in cities with different ratio of CNR workers to total population. Bins refer to deviation from population weighted average.

each city/industry to 2011-15 levels while keeping the relative productivity of CNR and non-CNR workers at 1980 levels. Relative to the baseline (rows 4 and 15), CNR workers are worse off and non-CNR workers better off. The difference is accounted for skill-biased technical change. Note that externalities amplify the effect of skill biased technical change.

Rows 5 and 16 change the composition of employment to 2011-15 levels, with more CNR workers and fewer non-CNR workers. Here, externalities play their largest role. Absent externalities, the model would imply significant losses for CNR workers, as they become more abundant, whereas standard neo-classical arguments would imply a reduction in their relative wage. In effect, absent externalities, CNR workers would end up with welfare 14% below the baseline counterfactual, while non-CNR worker's welfare would grow by 11%. In contrast, with externalities, both CNR and non-CNR workers end up gaining about the same as occupations become more polarized across space.

Rows 6 and 17 add exogenous changes to local technology of non-real estate sectors over and above what is implied by average national trends. It captures, for example, the fact that computer and electronics output became particularly more productive in San Jose while finance became particularly more productive in New York. As shown in Figure 18 those gains were larger in cities that had high CNR shares in 1980. These location-specific technological changes interact with externalities to increase the welfare of both types of workers.

Rows 7 and 18 add changes in the productivity of the real estate sector. This exercise encompasses the effects of two different underlying processes. On the one hand, real estate productivity increased more in fast growing cities, as the stock of housing increased in order to accommodate rising populations. On the other hand, as it has been increasingly recognized (Glaeser and Gyourko (2018), Hsieh and Moretti (2019) and Herkenhoff et al. (2018)), housing regulations have impeded

Table 8: Welfare Comparison, Relative to Baseline

	CNR	non-CNR	CNR-to-non-CNR Ratio
Full Model			
1. 1980 parameters	0.558	0.765	0.730
2. (1) + current input shares	0.708	0.824	0.859
3. (2) + national trends in technology and population	0.927	1.081	0.857
4. (3) + national skill biased technical change (Baseline)	1.000	1.000	1.000
ratio of CNR to non-CNR welfare in Baseline		1.911	
5. (4) + change in occ. shares in employment	1.045	1.058	0.987
6. (5) + change in local technology (ex real estate)	1.066	1.079	0.988
7. (6) + change in real estate productivity	1.042	1.060	0.983
8. (7) + change in amenities (2011-15 parameters)	1.039	1.061	0.980
9. Optimal Allocation	1.045	1.067	0.980
10. 2011-15 parameters minus change in real estate productivity	1.054	1.070	0.985
11. Optimal Allocation with parameters in (10)	1.064	1.081	0.985
Model without externalities			
12. 1980 parameters	0.488	0.643	0.759
13. (12) + current input shares	0.790	0.893	0.885
14. (13) + national trends in technology and population	0.936	1.057	0.885
15. (14) + national skill biased technical change (Baseline)	1.000	1.000	1.000
ratio of CNR to non-CNR welfare in Baseline		3.767	
16. (15) + change in occ. shares in employment	0.858	1.110	0.773
17. (16) + change in local technology (ex real estate)	0.852	1.116	0.763
18. (17) + change in real estate productivity	0.860	1.129	0.762
19. (18) + change in amenities (2011-15 parameters)	0.861	1.129	0.762
20. Optimal Allocation	0.863	1.132	0.762
21. 2011-15 parameters minus change in real estate productivity	0.846	1.109	0.763
22. Optimal Allocation with parameters in (21)	0.849	1.112	0.763

development in some very productive areas. The table shows that the net effect of those two forces would have been positive in the absence of externalities, as housing development may have accommodated increasing population in high growing locations. At the same time, their effect is negative once externalities are accounted for since, as shown in Figure 19, housing productivity also lagged behind in CNR intensive cities.

Finally, Rows 8 and 19 add the changes in amenities. In particular, Row 8 corresponds to the 2011-15 equilibrium allocation. Changes in the spatial distribution of amenities appear to add little to total welfare. That said, in considering the results presented in Table 8, it is important to bear in mind that the particular sequence in which we added the changes between 1980 and 2011-15 can have an effect on our results. We chose to present a sequence that is intuitive to us, but the main findings highlighted above are robust to other sequences.

We calculate optimal allocation under two scenarios. The first scenario, depicted in Rows 9 and 20, corresponds to that depicted in Section 5 above when externalities are included. The second scenario assesses the role of housing policy in impeding optimal policy. In particular, we calculate the optimal policy under the assumption that housing productivity was distributed as in 1980 (the corresponding equilibrium counterfactual welfare is presented in Rows 10 and 21 and the corresponding optimal allocation in Rows 11 and 22). We find that the increment in welfare is more than 50% larger in that scenario than when starting from the actual equilibrium. In other words, the optimal policy is less effective in some of the cognitive hubs owing to observed changes in housing supply restrictions .

7 Conclusion

Our aim in this paper has been to understand the extent to which workers are misallocated in space and the policies that might improve observed allocations. The main culprit of spatial misallocation is the existence of large occupation-specific externalities combined with potential distortions due to land use regulations. Our quantitative spatial model allowed us to measure occupation-specific local productivity by industry which, together with a relatively standard instrumental variable approach, led us to estimate these externalities for CNR and non-CNR occupations.

Our estimates suggest that both CNR and non-CNR workers become more productive in large cities, but CNR productivity improves particularly when CNR workers are surrounded by other CNR workers. These estimates, together with estimated local amenities by occupation, exogenous productivity differences across industries and locations, and the full set of input-output linkages

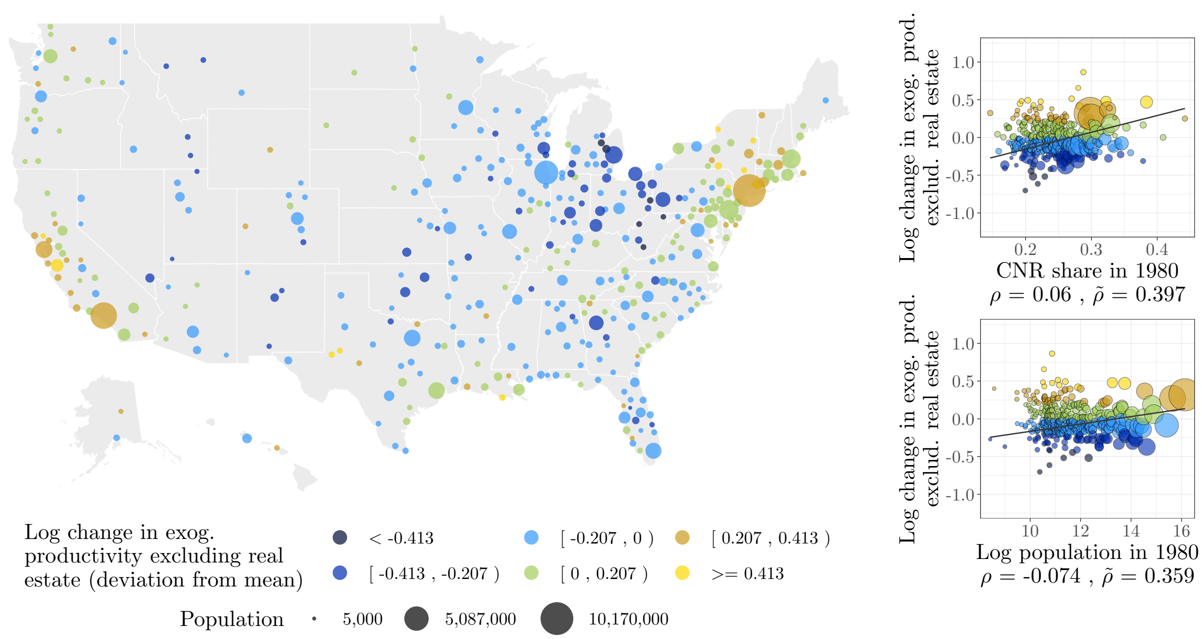

Figure 18: Change in the exogenous part of technology (all non-real estate sectors)

Changes are relative to baseline counterfactual, averaged across all sectors except real estate. Each observation refers to a CBSA. Marker sizes are proportional to total city employment. Averages are taken with value added weights. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

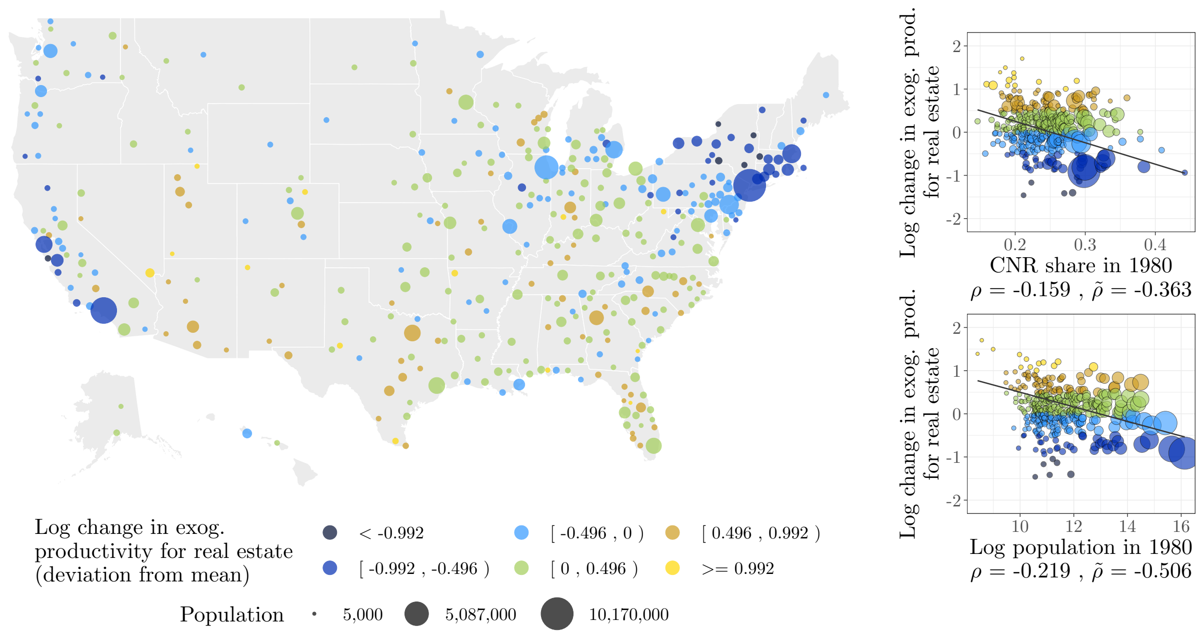

Figure 19: Change in the exogenous component of technology in real estate.

Changes are relative to baseline counterfactual. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

and transport costs in the U.S. economy, determine the current allocation of economic activity. We find that an optimal spatial policy can improve on this allocation for both occupations by 0.59%. Housing and optimal transfer policies reinforce each other. Hence, combining them (by reverting the spatial distribution of real estate productivity to that of 1980) leads to welfare gains of close to 2.4% for CNR workers and 2% for non-CNRs.

Since the 80's the U.S. economy has experienced increased skill and occupational polarization across space. Large cities increasingly have more highly educated CNR workers that earn more. In contrast, many medium and small cities have suffered an exodus of skilled workers and experienced persistent population declines. These trends, amplified by local externalities, were also associated with a rise in income inequality between occupations. This growing gap between top and medium and small-sized cities has motivated policymakers and city governments to advocate policies to attract CNR workers to smaller towns in order to reverse their fortunes. Our analysis shows that, given appropriate transfers, these efforts would be counterproductive.

Our analysis underscores that while CNR workers are extremely useful, they are also scarce. Furthermore, their productivity is tremendously enhanced by living with other CNR workers. So attracting them to smaller towns with more mixed populations represents a waste of resources. CNR workers are too valuable for society to be used in this way. A better policy is to reinforce existing trends and let them concentrate in cognitive hubs while incentivizing non-CNR workers to move and help smaller cities grow. Of course, some non-CNR workers will always be needed in those hubs because of imperfect substitutability of occupations in production. The result is smaller, more CNR intensive, cognitive hubs in some of today's largest cities. We show that the resulting migration of non-CNR workers that allows small towns to grow may be implemented with a baseline transfer to non-CNR workers, reminiscent of a universal basic income, and a set of occupation-location specific transfers. Overall, CNR workers transfer resources to non-CNR workers to generate equal welfare gains.

Our findings suggest that efforts to stop the spatial polarization of occupations are misguided. In fact, encouraging it further can yield benefits for everyone when accompanied by the necessary transfers. Implementing these transfers, however, is critical. Otherwise, cognitive hubs might use other indirect means of pushing out non-CNR workers such as, for example, housing supply constraints, zoning restrictions, or a lack of investment in transportation networks to aid commuting. Such efforts can generate occupational polarization across space without Pareto gains for all workers. Implementing the necessary transfers would not only help avoid those inefficient policies and benefit CNR workers, but it would also improve the welfare of non-CNR workers and the many small and medium sized cities where they would end up living, working, and producing.

Our analysis abstracts from the role that spatial polarization might have on human capital formation. In principle, the migration of CNR workers towards cognitive hubs may be detrimental to smaller cities in a setting where the learning technology also features meaningful externalities from CNR workers. At the same time, however, the transfers that CNR workers are able and willing to make to non-CNR workers, given the productivity gains they experience from living in cognitive hubs, might naturally be invested in education and other training in the smaller cities. These transfers, if directed properly, have the potential to ameliorate, or even reverse, the conceivably negative effects of spatial polarization on human capital.

References

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier. 2

Albouy, D. (2009). The unequal geographic burden of federal taxation. *Journal of Political Economy 117*(4), 635–667. 1

Albouy, D., K. Behrens, F. Robert-Nicoud, and N. Seegert (2019). The optimal distribution of population across cities. *Journal of Urban Economics 110*, 102–113. 1

Altonji, J. G., T. E. Elder, and C. R. Taber (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy 113*(1), 151–184. B.3

Anderson, J. E., C. A. Milot, and Y. V. Yotov (2014). How much does geography deflect services trade? canadian answers. *International Economic Review 55*(3), 791–818. 4, B.1, 9

Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion.* Princeton university press. 26

Autor, D. H. (2019). Work of the past, work of the future. In *Richard T. Ely Lecture given at Annual Meeting of American Economics Association, Atlanta, GA, January*, Volume 4. 1

Bartelme, D., A. Costinot, D. Donaldson, and R.-C. Andres (2019). The textbook case for industrial policy: Theory meets data. Technical report, CEPR Discussion Papers. 5.3.1

Baum-Snow, N., M. Freedman, and R. Pavan (2018). Why has urban inequality increased? *American Economic Journal: Applied Economics 10*(4), 1–42. 1

Baum-Snow, N. and R. Pavan (2011). Understanding the city size wage gap. *The Review of economic studies 79*(1), 88–127. 9

Baum-Snow, N. and R. Pavan (2013). Inequality and city size. *Review of Economics and Statistics 95*(5), 1535–1548. 1, 2, 16

Berry, C. R. and E. L. Glaeser (2005). The divergence of human capital levels across cities. *Papers in regional science 84*(3), 407–444. 1

Brinkman, J. (2014). The supply and demand of skilled workers in cities and the role of industry composition. 1

Burstein, A., E. Morales, and J. Vogel (2019). Changes in between-group inequality: computers, occupations, and international trade. *American Economic Journal: Macroeconomics 11*(2), 348–400. 10

Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies 82*(1), 1–44. 4, 13

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2017). The impact of regional and sectoral productivity changes on the us economy. *The Review of economic studies 85*(4), 2042–2096. 1, 4.1.3, A.2.3, A.6

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics 19*(1), 22–64. 1, 4, 4.2.1, B.3, B.3

Ciccone, A. and R. E. Hall (1996). Productivity and the density of economic activity. *The American Economic Review 86*(1), 54. 4.2.1, B.3, B.3

Ciccone, A. and G. Peri (2005). Long-run substitutability between more and less educated workers: evidence from us states, 1950–1990. *Review of Economics and statistics 87*(4), 652–663. 4

Colas, M. and K. Hutchinson (2017). Heterogeneous workers and federal income taxes in a spatial equilibrium. Technical report. 1

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2018). Income growth and the distributional effects of urban spatial sorting. Technical report, Working Paper University of Chicago Booth. 4

Davis, M. A. and J. Gregory (2020). Place-based redistribution in simple location-choice models. 30

Desmet, K. and E. Rossi-Hansberg (2013). Urban accounting and welfare. *American Economic Review 103*(6), 2296–2327. 1

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review 106*(3), 479–524. 1, 2, 1, 4.1.2, 18, 5.3.1, E

Duranton, G. and D. Puga (2004). Micro-foundations of urban agglomeration economies. In *Handbook of regional and urban economics*, Volume 4, pp. 2063–2117. Elsevier. 3

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779. 3.2.2, A.2.3

Eeckhout, J. and N. Guner (2015). Optimal spatial taxation: Are big cities too small? 1

Ellison, G., E. L. Glaeser, and W. R. Kerr (2010). What causes industry agglomeration? evidence from coagglomeration patterns. *American Economic Review 100*(3), 1195–1213. 1, 4.2

Fajgelbaum, P. and C. Gaubert (2018). Optimal spatial policies, geography and sorting. Technical report, National Bureau of Economic Research. 1, 1, 18, 5.1, 30, E, 21

Fajgelbaum, P. D., E. Morales, J. C. Suárez Serrato, and O. Zidar (2018). State taxes and spatial misallocation. *The Review of Economic Studies 86*(1), 333–376. 1, 4, 15

Feenstra, R. C., P. Luck, M. Obstfeld, and K. N. Russ (2018). In search of the armington elasticity. *Review of Economics and Statistics 100*(1), 135–150. 13

Giannone, E. (2017). Skilled-biased technical change and regional convergence. 1

Glaeser, E. and J. Gyourko (2018, February). The economic implications of housing supply. *Journal of Economic Perspectives 32*(1), 3–30. 1, 6

Glaeser, E. L. and J. D. Gottlieb (2008). The economics of place-making policies. *Brookings Papers on Economic Activity 2008*(1), 155–253. 21

Handbury, J. and D. E. Weinstein (2014). Goods prices and availability in cities. *The Review of Economic Studies 82*(1), 258–296. 4.1.1

Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of international economics*, Volume 4, pp. 131–195. Elsevier. 4

Hendricks, L. (2011). The skill composition of us cities. *International Economic Review 52*(1), 1–32. 1

Herkenhoff, K. F., L. E. Ohanian, and E. C. Prescott (2018). Tarnishing the golden and empire states: Land-use restrictions and the us economic slowdown. *Journal of Monetary Economics 93*, 89–109. 1, 1, 6

Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*. 1, 1, 6

Jaimovich, N. and H. E. Siu (2018). The trend is the cycle: Job polarization and jobless recoveries. *The Review of Economics and Statistics*. 5

Lucas, R. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics 22*, 3–42. 1

Marshall, A. (1920). *Principles of Economics (Revised ed.)*. London: Macmillan; reprinted by Prometheus Books. 1

Melo, P. C., D. J. Graham, and R. B. Noland (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional science and urban Economics 39*(3), 332–342. 1, 4.1.3, 4

Moretti, E. (2004a). Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data. *Journal of econometrics 121*(1-2), 175–212. 1, 1, 16, 1, 4.2.1, B.3, B.3

Moretti, E. (2004b). Workers' education, spillovers, and productivity: evidence from plant-level production functions. *American Economic Review 94*(3), 656–690. 1, 4.1.3, 4

Neal, D. (1997). The effects of catholic secondary schooling on educational achievement. *Journal of Labor Economics 15*(1, Part 1), 98–123. B.3

Neumark, D. and H. Simpson (2015). Place-based policies. In *Handbook of regional and urban economics*, Volume 5, pp. 1197–1287. Elsevier. 1

Ossa, R. (2015). A quantitative analysis of subsidy competition in the us. Technical report, National Bureau of Economic Research. 1

Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics 9*, 21–58. 1

Roca, J. D. L. and D. Puga (2017). Learning by working in big cities. *The Review of Economic Studies 84*(1), 106–142. 9, 1

Sanderson, E. and F. Windmeijer (2016). A weak instrument f-test in linear iv models with multiple endogenous variables. *Journal of Econometrics 190*(2), 212–221. 4.2.1

Staiger, D., J. H. Stock, et al. (1997). Instrumental variables regression with weak instruments. *Econometrica 65*(3), 557–586. 4.2.1

Stock, J. H. and M. Yogo (2005). Testing for weak instruments in linear iv regression. *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 80. 4.2.1

Appendix for online publication

A Model Details

A.1 Household Decisions

In a given occupation, all households living in the same city choose the same consumption basket. It follows that $C_n^{kj}(\mathbf{a}) = C_n^{kj}$ for all $\mathbf{a}$. Moreover, the demand for good j by workers in occupation k living in city n is

$$C_n^{kj} = \alpha^j \frac{P_n}{P_n^j} C_n^k \tag{15}$$

where $P_n = \prod_{j=1}^{J} \left(\frac{P_n^j}{\alpha^j} \right)^{\alpha^j}$ is the ideal price index in city n.

Agents move freely across cities. The value, $v_n^k(\mathbf{a})$, of locating in a particular city n for an individual employed in occupation k, with idiosyncratic preference vector $\mathbf{a}$ is

$$v_n^k(\mathbf{a}) = \frac{a_n A_n^k I_n^k}{P_n} = a_n A_n^k C_n^k.$$

In equilibrium, workers move to the location where they receive the highest utility so that

$$v^k(\mathbf{a}) = \max_n v_n^k(\mathbf{a}),$$

where $v^k(\mathbf{a})$ now denotes the equilibrium utility of an individual in occupation k with amenity vector $\mathbf{a}$. We assume that a_n is drawn from a Fréchet distribution. Draws are independent across cities. We denote by Ψ the joint *cdf* for the elements of $\mathbf{a}$ across workers in a given occupation, with

$$\Psi(\mathbf{a}) = \exp\left\{ -\sum_n (a_n)^{-\nu} \right\},$$

where the shape parameter ν reflects the extent of preference heterogeneity across workers. Higher values of ν imply less heterogeneity, with all workers ordering cities in the same way when $\nu \to \infty$.

Assuming that workers of different types can freely move between cities, the average utility of a worker of type k is given by

$$v^k = \Gamma\left(\frac{\nu - 1}{\nu}\right) \left(\sum_n \left(A_n^k C_n^k \right)^{\nu} \right)^{\frac{1}{\nu}}, \tag{16}$$

where $\Gamma(.)$ is the Gamma function.

Combining this equation with equation (2) describing labor supply yields an expression relating the value of each occupational type to consumption and employment in particular locations:

$$v^k = \left(\frac{L_n^k}{L^k} \right)^{-\frac{1}{\nu}} A_n^k C_n^k.$$

A.2 Firms

A.2.1 Intermediate Goods Producers

Cost minimization implies that input demand satisfies:

$$\frac{r_n H_n^j(\mathbf{z})}{x_n^j(\mathbf{z}) q_n^j(\mathbf{z})} = \gamma_n^j \beta_n^j, \tag{17}$$

$$\frac{w_n^k L_n^{kj}(\mathbf{z})}{x_n^j(\mathbf{z}) q_n^j(\mathbf{z})} = \frac{\left(\frac{w_n^k}{\lambda_n^{kj}}\right)^{1-\epsilon}}{\sum_{k'}\left(\frac{w_n^{k'}}{\lambda_n^{k'j}}\right)^{1-\epsilon}} \gamma_n^j \left(1-\beta_n^j\right), \tag{18}$$

$$\frac{P_n^{j'} M_n^{j'j}(\mathbf{z})}{x_n^j(\mathbf{z}) q_n^j(\mathbf{z})} = \gamma_n^{j'j}, \tag{19}$$

where $x_n^j(\mathbf{z})$ is the Lagrange multiplier which in this case reflects the unit cost of production. We can solve for $x_n^j(\mathbf{z})$ by substituting optimal factor choices into the production function,

$$x_n^j(\mathbf{z}) \equiv \frac{x_n^j}{z_n} = \frac{B_n^j}{z_n}\left\{\frac{r_n^{\beta_n^j}}{Z_n^j}\left[\sum_k\left(\frac{w_n^k}{\lambda_n^{kj}}\right)^{1-\epsilon}\right]^{\frac{1-\beta_n^j}{1-\epsilon}}\right\}^{\gamma_n^j} \prod_{j'=1}^{J}\left(P_n^{j'}\right)^{\gamma_n^{j'j}} \tag{20}$$

where x_n^j is a city and industry specific unit cost index such that

$$B_n^j = \left[\left(1-\beta_n^j\right)^{\beta_n^j - 1}\left(\beta_n^j\right)^{-\beta_n^j}\right]^{\gamma_n^j}\left[\prod_{j'}\left(\gamma_n^{j'j}\right)^{-\gamma_n^{j'j}}\right]\left(\gamma_n^j\right)^{-\gamma_n^j}.$$

Given constant returns to scale and competitive intermediate goods markets, a firm produces positive but finite amounts of a variety only if its price is equal to its unit production cost,

$$p_n^j(\mathbf{z}) = x_n^j(\mathbf{z}) = \frac{x_n^j}{z_n}. \tag{21}$$

A.2.2 Final Goods

Let $Q_n^j(\mathbf{z}) = \sum_{n'} Q_{nn'}^j(\mathbf{z})$ denote the total amount of intermediate goods of variety $\mathbf{z}$ purchased from different cities by a final goods producer in city n, sector j. Given that intermediate goods of a given variety produced in different cities are perfect substitutes, final goods producers purchase varieties only from cities that offer the lowest unit cost,

$$Q_{nn'}^j(\mathbf{z}) = \begin{cases} Q_n^j(\mathbf{z}) & if\ \kappa_{nn'}^j p_{n'}^j(\mathbf{z}) < \min_{n''\neq n'} \kappa_{nn''}^j p_{n''}^j(\mathbf{z}) \\ 0 & \text{otherwise} \end{cases},$$

where we abstract from the case where $\kappa_{nn'}^j p_{n'}^j(\mathbf{z}) = \min_{n''\neq n'} \kappa_{nn''}^j p_{n''}^j(\mathbf{z})$ since, given the distributional assumption on $\mathbf{z}$, this event only occurs on a set of measure zero.

Denote by $P_n^j(\mathbf{z})$ the unit cost paid by a final good producer in city n and sector j for a particular variety whose vector of productivity draws is $\mathbf{z}$. Given that final goods firms only purchase intermediate goods from the lowest cost supplier,

$$P_n^j(\mathbf{z}) = \min_{n'} \left\{ \kappa_{nn'}^j p_{n'}^j(\mathbf{z}) \right\} = \min_{n'} \left\{ \frac{\kappa_{nn'}^j x_{n'}^j}{z_{n'}} \right\}. \tag{22}$$

For non-tradable intermediate goods, firms must buy those goods locally, so that if j is non-tradable,

$$P_n^j(\mathbf{z}) = \frac{x_n^j}{z_n}. \tag{23}$$

Then, the demand function for intermediate goods of variety $\mathbf{z}$ in industry j and city n is given by

$$Q_n^j(\mathbf{z}) = \left(\frac{P_n^j(\mathbf{z})}{\tilde{P}_n^j} \right)^{-\eta} Q_n^j, \tag{24}$$

where $\tilde{P}_n^j$ the ideal cost index for final goods produced in sector j in city n,

$$\tilde{P}_n^j = \left[\int P_n^j(\mathbf{z})^{1-\eta} d\Phi(\mathbf{z}) \right]^{\frac{1}{1-\eta}}. \tag{25}$$

Since the production function for final goods is constant returns to scale, and the market for final goods is competitive, a final goods firm produces positive but finite quantities of a final good if its price is equal to its cost index, that is if $P_n^j = \tilde{P}_n^j$.

A.2.3 Derivation of Prices

We follow Eaton and Kortum (2002) in solving for the distribution of prices. Given this distribution and zero profits for final goods producers, when sector j is tradable, the price of final goods in sector j in region n solves

$$\left(P_n^j\right)^{1-\eta} = \int P_n^j(\mathbf{z})^{1-\eta} d\Phi(\mathbf{z}) \, d\mathbf{z},$$

which is the expected value of the random variable $P_n^j(\mathbf{z})^{1-\eta}$.

Let $P_{nn'}^j(\mathbf{z}) = \frac{\kappa_{nn'}^j x_{n'}^j}{z_{n'}}$ denote the unit cost of a variety indexed by $\mathbf{z}$ produced in city n' and sold in n. Following the steps described in Caliendo et al. (2017), we have that

$$\Pr\left[P_{nn'}^j(\mathbf{z}) \leq p\right] = 1 - e^{-\omega_{nn'}^j p^\theta}$$

where $\omega_{nn'}^j = \left[\kappa_{nn'}^j x_{n'}^j\right]^{-\theta_j}$. The price of variety $\mathbf{z}$ in city n and industry j, $P_n^j(\mathbf{z})$, is the minimum across $P_{nn'}^j(\mathbf{z})$. Its *cdf* is,

$$\Pr\left[P_n^j(\mathbf{z}) \leq p\right] = 1 - e^{-\Omega_n^j p^\theta},$$

where $\Omega_n^j = \sum_{n'} \omega_{nn'}^j = \sum_{n'} \left[\kappa_{nn'}^j x_{n'}^j\right]^{-\theta}$ (Ω_n^j does not depend on n' because we are integrating out the city dimension).

Let $F_{P_n^j}(p)$ denote the *cdf* of $P_n^j(\mathbf{z})$, $\Pr\left[P_n^j(\mathbf{z}) \leq p\right]$. Then, its associated *pdf*, denoted $f_{P_n^j}(p)$, is $\Omega_n^j \theta p^{\theta-1} e^{-\Omega_n^j p^\theta}$. As in Caliendo et al. (2017), we have that

$$P_n^j = \Gamma(\xi)^{\frac{1}{1-\eta}} \left(\Omega_n^j\right)^{-\frac{1}{\theta}},$$

where $\Gamma(\xi)$ is the Gamma function evaluated at $\xi = 1 + \frac{1-\eta}{\theta}$. The price of goods in tradable sector j may then also be expressed as

$$P_n^j = \Gamma(\xi)^{\frac{1}{1-\eta}} \left[\sum_{n'=1}^{N} \left[\kappa_{nn'}^j x_{n'}^j \right]^{-\theta} \right]^{-\frac{1}{\theta}}.$$

In a given non-tradable sector j, $\kappa_{nn'}^j = \infty \; \forall n' \neq n$, so that the equation reduces to

$$P_n^j = \Gamma(\xi)^{\frac{1}{1-\eta}} x_n^j.$$

A.2.4 Trade Shares

Let X_n^j denote *total expenditures* on final goods j by city n, which must equal of the value of final goods in that sector, $X_n^j = P_n^j Q_n^j$. Recall that because of zero profits in the final goods sector, total expenditures on intermediate goods in a given sector are then also equal to the cost of inputs used in that sector, so that $P_n^j Q_n^j = \int P_n^j(\mathbf{z}) Q_n^j(\mathbf{z}) d\Phi(\mathbf{z})$. Let $X_{nn'}^j = \int \kappa_{nn'}^j p_{n'}^j(\mathbf{z}) Q_{nn'}^j(\mathbf{z}) d\Phi(\mathbf{z})$ denote the value spent by city n on intermediate goods of sector j produced in city n'. Further, let $\pi_{nn'}^j$ denote the share of city n's expenditures on sector j goods purchased from region n'. Then,

$$\pi_{nn'}^j = \frac{X_{nn'}^j}{X_n^j}.$$

Observe that, since there is a continuum of varieties of intermediate goods, the fraction of goods that firms in city n purchase from firms in city n' is given by

$$\tilde{\pi}_{nn'}^j \equiv \Pr\left[P_{nn'}^j(\mathbf{z}) \leq \min_{n'' \neq n'} \left\{ P_{nn''}^j(\mathbf{z}) \right\} \right].$$

Following the steps described in Caliendo et al. (2017), we have that

$$\begin{aligned} \tilde{\pi}_{nn'}^j &= \frac{\omega_{nn'}^j}{\Omega_n^j} \\ &= \frac{\left[\kappa_{nn'}^j x_{n'}^j\right]^{-\theta}}{\sum_{n''=1}^{N} \left[\kappa_{nn''}^j x_{n''}^j\right]^{-\theta}} \end{aligned}$$

We can verify that $\tilde{\pi}_{nn'}^j = \pi_{nn'}^j$, that is, the share of goods that firms in city n purchase from city n' is equal to the share of the *value* of goods produced in city n' in the bundle purchased by firms in city n (see Eaton and Kortum (2002), Footnote 17). Observe also that $\sum_{n'=1}^{N} \left[\kappa_{nn'}^j x_{n'}^j\right]^{-\theta} = \left(P_n^j\right)^{-\theta} \Gamma(\xi)^{\frac{\eta}{1-\eta}}$. Therefore, we may alternatively write the trade share $\pi_{nn'}^j$ as

$$\pi_{nn'}^j = \frac{X_{nn'}^j}{X_n^j} = \left[\frac{\kappa_{nn'}^j x_{n'}^j \Gamma(\xi)^{\frac{1}{1-\eta}}}{P_n^j} \right]^{-\theta}$$

In non-tradable sectors, $\pi_{nn}^j = 1$.

A.3 Market Clearing and Aggregation at the Industry and City Level

Given the labor supply equation (2) and the definition $L_n^{kj} = \int L_n^{kj}(\mathbf{z})d\Phi(\mathbf{z})$, the labor market clearing equation (5) may be rewritten as

$$\sum_j L_n^{kj} = L^k \frac{\left(A_n^k C_n^k\right)^\nu}{\sum_{n'} \left(A_{n'}^k C_{n'}^k\right)^{\nu^k}}, \; \forall \; n = 1, ..., N, \; k = 1, ..., K.$$

Given the definition $H_n^j = \int H_n^j(\mathbf{z})d\Phi(\mathbf{z})$, the market clearing equation for structures in each city (6) may be rewritten as

$$\sum_j H_n^j = H_n, \; n = 1, ..., N.$$

Given our definition of total final expenditures, $X_n^j = P_n^j Q_n^j$, and the demand function for consumption goods of sector j (15), the market clearing condition for final goods in each city n and sector j (7) may be expressed in terms of sectoral and city aggregates,

$$\sum_k L_n^k \left(\alpha^j P_n C_n^k\right) + P_n^j \sum_{j'} M_n^{jj'} = X_n^j.$$

Finally, given that $\pi_{n'n}^j X_{n'}^j = X_{n'n}^j = \int p_n^j(\mathbf{z}) \kappa_{n'n}^j Q_{n'n}^j(\mathbf{z}) d\Phi(\mathbf{z})$, the market clearing condition for intermediate inputs (8) may be rewritten in terms of sectoral city aggregates as

$$\underbrace{\int p_n^j(\mathbf{z}) q_n^j(\mathbf{z}) d\Phi(\mathbf{z})}_{\text{Total value of intermediate goods produced in city } n} = \sum_{n'} \pi_{n'n}^j X_{n'}^j,$$

where $\sum_{n'} \pi_{n'n}^j X_{n'}^j$ is the total value of expenditures across all cities spent on intermediate goods produced in city n.

We can use this last aggregation relationship to obtain aggregate factor input demand equations as follows,

$$\begin{aligned}
w_n^k L_n^{kj} &= \gamma_n^j \left(1 - \beta_n^j\right) \frac{\left(\frac{w_n^k}{\lambda_n^{kj}}\right)^{1-\epsilon}}{\sum_{k'=1}^K \left(\frac{w_n^{k'}}{\lambda_n^{k'j}}\right)^{1-\epsilon}} \sum_{n'} \left(\pi_{n'n}^j X_{n'}^j\right), \\
r_n H_n^j &= \gamma_n^j \beta_n^j \sum_{n'} \left(\pi_{n'n}^j X_{n'}^j\right), \\
P_n^{j'} M_n^{jj'} &= \gamma_n^{j'j} \sum_{n'} \left(\pi_{n'n}^j X_{n'}^j\right).
\end{aligned}$$

Finally, combining these factor demand equations yields the aggregate production function,

$$\sum_{n'} \pi_{n'n}^j X_{n'}^j = x_n^j \left[\left(\sum_k \left(\lambda_n^{kj} L_n^{kj}\right)^{1-\frac{1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}(1-\beta_n^j)} \left(H_n^j\right)^{\beta_n^j} \right]^{\gamma_n^j} \prod_{j'} \left(M_n^{j'j}\right)^{\gamma_n^{j'j}}.$$

A.4 Definition of Equilibrium

Equilibrium for this system of cities is given by a set of final goods prices P_n^j, wages in different occupations, w_n^k, rental rates, r_n, intermediate goods prices paid by final goods producers, $P_n^j(\mathbf{z})$, intermediate goods prices received by intermediate goods producers, $p_n^j(\mathbf{z})$, consumption choices, C_n^{kj}, intermediate input choices, $Q_n^j(\mathbf{z})$, intermediate input production, $q_n^j(\mathbf{z})$, demand for materials, $M_n^{jj'}(\mathbf{z})$, labor demand, $L_n^{kj}(\mathbf{z})$, demand for structures, $H_n(\mathbf{z})$, and location decisions, $\zeta_n^k(\mathbf{a})$, such that:

i) Workers choose consumption of each final good optimally, as implied by equation (15) and the budget constraint, $\sum_j P_n^j C_n^{kj} = P_n C_n^k = I_n^k$, where $P_n = \prod_j \left(\frac{P_n^j}{\alpha^j}\right)^{\alpha^j}$ and I_n^k is given by equation (1).

ii) Workers choose optimally where to live as implied by equation (2).

iii) Intermediate input producers choose their demand for materials, labor and structures optimally (as implied by factor demand equations (17), (18) and (19)), and produce positive but finite amounts only if (21) holds, where x_n^j in that equation is given by (20).

iv) Final goods producers choose the origin of intermediate inputs optimally, implying that a producer in city n and industry j imports a variety $\mathbf{z}$ from city n' if and only if $\kappa_{nn'}^j p_n^j(\mathbf{z}) = \min_{n''}\left\{\kappa_{nn''}^j p_{n''}^j(\mathbf{z})\right\}$. The price that they pay for intermediate goods satisfies (22) if the good is tradable and (23) if it is non-tradable.

v) Final goods producers choose their intermediate input use optimally according to (24) and produce positive but finite amounts only if (25) holds.

vi) Market clearing conditions for employment (equation 5), land and structures (equation 6), final goods (equation 7), and intermediate goods (equation 8) hold.

A.5 Aggregate Equilibrium

At the aggregate level, equilibrium is given by values for the prices P_n, P_n^j, x_n^j, r_n, w_n^k, aggregate quantities C_n^k, L_n^{kj}, H_n^j, $M_n^{j'j}$, expenditures, X_n^j, and expenditure shares, $\pi_{nn'}^j$, that satisfy the following equations

$$\sum_{k,j'} L_n^{kj'}\left(\alpha^j P_n C_n^k\right) + \sum_{j'} P_n^j M_n^{jj'} = X_n^j \quad (NJ \text{ eqs.}) \tag{26}$$

$$L_n^k = \sum_j L_n^{kj} = \frac{\left(A_n^k C_n^k\right)^\nu}{\sum_{n'}\left(A_{n'}^k C_{n'}^k\right)^\nu} L^k \quad (NK \text{ eqs.}) \tag{27}$$

$$\sum_j H_n^j = H_n \quad (N \text{ eqs.}) \tag{28}$$

$$P_n = \prod_j \left(\frac{P_n^j}{\alpha^j}\right)^{\alpha^j} \quad (N \text{ eqs.}) \tag{29}$$

$$P_n C_n^k = w_n^k + b^k \frac{\sum_{n'} r_{n'} H_{n'}}{L^k} \text{ where } b^k = \frac{\sum_n w_n^k L_n^k}{\sum_{n,k'} w_n^{k'} L_n^{k'}} \; (NK \text{ eqs.}) \tag{30}$$

$$w_n^k L_n^{kj} = \frac{\left(\frac{w_n^k}{\lambda_n^{kj}(\mathbf{L}_n)}\right)^{1-\epsilon}}{\sum_{k'}\left(\frac{w_n^{k'}}{\lambda_n^{k'j}(\mathbf{L}_n)}\right)^{1-\epsilon}} \gamma_n^j \left(1-\beta_n^j\right) \sum_{n'} \pi_{n'n}^j X_{n'}^j \; (NKJ \text{ eqs.}) \tag{31}$$

$$r_n H_n^j = \gamma_n^j \beta_n^j \sum_{n'} \pi_{n'n}^j X_{n'}^j \ \ (NJ \text{ eqs.}) \tag{32}$$

$$P_n^{j'} M_n^{j'j} = \gamma_n^{j'j} \sum_{n'} \pi_{n'n}^j X_{n'}^j \ \ \left(NJ^2 \text{ eqs.}\right) \tag{33}$$

$$P_n^j = \begin{cases} \Gamma(\xi)^{\frac{1}{1-\eta}} \left(\sum_{n'} \left[\kappa_{nn'}^j x_{n'}^j \right]^{-\theta} \right)^{-\frac{1}{\theta}} & \text{if } j \text{ is tradable} \\ \Gamma(\xi)^{\frac{1}{1-\eta}} x_n^j & \text{if } j \text{ is non-tradable} \end{cases} \quad (NJ \text{ eqs.}) \tag{34}$$

$$\begin{aligned} & \sum_{n'} \pi_{n'n}^j X_{n'}^j \\ = \ \ & x_n^j \left[\left(\sum_k \left(\lambda_n^{kj}(\mathbf{L}_n) L_n^{kj} \right)^{1-\frac{1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}(1-\beta_n^j)} \left(H_n^j\right)^{\beta_n^j} \right]^{\gamma_n^j} \prod_{j'} \left(M_n^{j'j}\right)^{\gamma_n^{j'j}} \ \ (NJ \text{ eqs.}) \end{aligned} \tag{35}$$

$$\pi_{nn'}^j = \frac{\left[\kappa_{nn'}^j x_{n'}^j\right]^{-\theta}}{\sum_{n''} \left[\kappa_{nn''}^j x_{n''}^j\right]^{-\theta}} \ \left(N^2 J \text{ eqs.}\right) \tag{36}$$

This system of equations comprises $2N + 2NK + 4NJ + NKJ + NJ^2 + N^2J$ equations in the same number of unknowns.

By substituting equation (33) into equation (26), adding over all industries (j) and all cities (n) and rearranging, we arrive at the National Accounting identity stating that aggregate value added is equal to aggregate consumption expenditures in the economy,

$$\sum_{n,k,j} L_n^{kj} P_n C_n^k = \sum_{n,j} \gamma_n^j X_n^j \tag{37}$$

At the same time, multiplying both sides of equation (30) by L_n^k, adding over city (n) and occupation (k), and substituting out w_n^k and $r_n H_n$ using equations (28), (31) and (32), yields the same national accounting identity. The fact that we can arrive at that same identity by manipulating different sets of equations implies that there is one redundant equation in the system, leading to one too many unknowns relative to the number of equations. The presence of a redundant equation is a feature of Walrasian systems. In order to pin down the price level, therefore, we need to amend the system with an additional equation defining the numeraire. Specifically, we set :

$$\sum_{n,j} \omega_n \ln(P_n) = \ln(\bar{P}), \tag{38}$$

where ω_n are a set of weights. When computing counterfactuals, we set those weights to be proportional to local nominal consumption: $\omega_n \propto P_n \sum_k C_n^k L_n^k$. Finally, observe that if we substitute the factor demand equations (31), (32), (33) into (35), we obtain the expression for the unit cost index,

$$x_n^j = B_n^j \left\{ r_n^{\beta_n^j} \left[\sum_{k=1}^K \left(\frac{w_n^k}{\lambda_n^{kj}} \right)^{1-\epsilon} \right]^{\frac{1-\beta_n^j}{1-\epsilon}} \right\}^{\gamma_n^j} \prod_{j'=1}^J \left(P_n^{j'}\right)^{\gamma_n^{j'j}}. \tag{39}$$

A.6 TFP accounting

In the text, we define TFP in equation (9) as

$$\ln TFP_n^j = \ln\left(\sum_{n'} \pi_{n'n}^j X_{n'}^j\right) - \ln P_n^j - \gamma_n^j \beta_n^j \ln H_n^j - \gamma_n^j (1-\beta_n^j) \sum_k \delta^{kj} \ln L_n^{kj} - \sum_{j'} \gamma_n^{j'j} d \ln M_n^{j'j},$$

where

$$\delta^{kj} = \frac{\sum_{n'} w_{n'}^k L_{n'}^{kj}}{\sum_{n,k'} w_{n'}^{k'} L_{n'}^{k'j}}$$

is the national share of occupation k in the wage bill across all occupations.

From equation (35,)

$$\sum_{n'} \pi_{n'n}^j X_{n'}^j = x_n^j \left[\left[\left(\sum \left(\lambda_n^{kj} L_n^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}}\right]^{1-\beta_n^j} \left(H_n^j\right)^{\beta_n^j}\right]^{\gamma_n^j} \prod_{j'} \left(M_n^{j'j}\right)^{\gamma_n^{j'j}}$$

Also, recall that we can write $\frac{x_n^j}{P_n^j} = \frac{1}{\kappa_{nn}^j}\left(\pi_{nn}^j\right)^{-\frac{1}{\theta}}$, which picks up the role of selection effects on productivity (see Caliendo et al. (2017)). We take a first order log-linear approximation of $\sum_{n'} \pi_{n'n}^j X_{n'}^j$ around national averages to obtain

$$d\ln\left(\sum_{n'} \pi_{n'n}^j X_{n'}^j\right) \simeq d\ln P_n^j - \frac{1}{\theta} d\ln \pi_{nn}^j + \gamma_n^j (1-\beta_n^j) \sum_k \frac{\left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}{\sum_{k'} \left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}} \left(d\ln L_n^{kj} + d\ln \lambda_n^{kj}\right)$$
$$+ \gamma_n^j \beta_n^j d\ln H_n^j + \sum_{j'} \gamma_n^{j'j} d\ln M_n^{j'j},$$

where $\frac{\left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}{\sum_{k'}\left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}$ is the national average of $\frac{\left(\lambda_n^{kj} L_n^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}{\sum_{k'}\left(\lambda_n^{kj} L_n^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}$. From manipulating equation (31), we can verify that:

$$\frac{\left(\lambda_n^{kj} L_n^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}{\sum_{k'}\left(\lambda_n^{kj} L_n^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}} = \frac{w_n^k L_n^{kj}}{\sum_{k'} w_n^{k'} L_n^{k'j}}$$

If we log-linearize around a national weighted average across cities, where we weight individual cities by their wage bill, we have that

$$\frac{\left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}}{\sum_{k'} \left(\lambda^{kj} L^{kj}\right)^{\frac{\epsilon-1}{\epsilon}}} = \sum_n \frac{w_n^k L_n^{kj}}{\sum_{k'} w_n^{k'} L_n^{k'j}} \times \frac{\sum_{k'} w_n^{k'} L_n^{k'j}}{\sum_{n',k'} w_{n'}^{k'} L_{n'}^{k'j}} = \frac{\sum_n w_n^k L_n^{kj}}{\sum_{n',k'} w_{n'}^{k'} L_{n'}^{k'j}} = \delta^{kj},$$

so that

$$d\ln\left(\sum_{n'}\pi^j_{n'n}X^j_{n'}\right)\simeq d\ln P^j_n-\frac{1}{\theta}d\ln\pi^j_{nn}+\gamma^j_n(1-\beta^j_n)\sum_k\delta^{kj}\left(d\ln L^{kj}_n+d\ln\lambda^{kj}_n\right)$$
$$+\gamma^j_n\beta^j_n d\ln H^j_n+\sum_{j'}\gamma^{j'j}_n d\ln M^{j'j}_n$$

Comparing to the expression for TFP, it follows that, up to a first order approximation,

$$d\ln TFP^j_n\simeq-\frac{1}{\theta}d\ln\pi^j_{nn}+\gamma^j_n(1-\beta^j_n)\sum_k\delta^{kj}\left(d\ln\lambda^{kj}_n\right)$$

which, abstracting from selection effects, reduces to

$$d\ln TFP^j_n=\gamma^j_n(1-\beta^j_n)\sum_k\delta^{kj}\left(d\ln\lambda^{kj}_n\right)$$

Defining $T^{kj}_n\equiv\left(\lambda^{kj}_n\right)^{(1-\beta^j_n)\gamma^j_n}\left(H^{kj}_n\right)^{\beta^j_n\gamma^j_n}$ it follows that for tradable sectors (in which case $\beta^j_n=0$ and $\gamma^j_n=\gamma^j$ for all n),

$$d\ln TFP^j_n=\sum_k\delta^{kj}d\ln T^{kj}_n$$

or

$$\ln TFP^j_n=\sum_k\delta^{kj}\ln T^{kj}_n+\text{constant independent of }n$$

For the purposes of comparing TFP^j_n across space, we can omit that constant.

B Quantifying the Model and Model Inversion

We now provide additional detail on how we quantify the model. The set of parameters needed to quantify our framework fall into broadly two types: i) parameters that are constant across cities (but may vary across occupations and/or industries) and that are directly available from statistical agencies, or that may be chosen to match national or citywide averages, and ii) parameters that vary at a more granular level and require using all of the model's equations (i.e. by way of model inversion) to match data that vary across cities, industries, and occupations.

B.1 Details on Parameters That Are Constant Across Cities

Input use shares in gross output (γ^j_n, $\gamma^{jj'}_n$) : To obtain an initial calibration for these share parameters, we use an average of the 2011 to 2015 BEA Use Tables, each adjusted by the same year's total gross output. The Use Table divides the value of the output in each sector j, $\sum_{n',n}\pi^j_{n'n}X^j_{n'}=\sum_n X^j_n$, into the value of input purchases from other sectors j', $\sum_{n,j'}P^{j'}_n M^{j'j}_n$, labor compensation, $\sum_{n,k}w^k_n L^{kj}_n$, operational surplus, $\sum_n r_n H^j_n$, and taxes on production and imports, $-\sum_n s^j_n X^j_n$,

$$\sum_n X_n^j = \sum_{n,j'} P_n^{j'} M_n^{j'j} + \sum_{n,k} w_n^k L_n^{kj} + \sum_n r_n H_n^j - \sum_n s_n^j X_n^j. \tag{40}$$

Input purchases from other sectors are separated into purchases from domestic producers and purchases from international producers. Since the model does not allow for foreign trade, we adjust the Use Table by deducting purchases from international producers from the input purchases and, for accounting consistency, from the definition of gross output for the sector.

As in Caliendo et al. (2018), for all sectors, we augment material purchases to include the purchases of equipment. Specifically, we subtract from the operational surplus of each sector 17 percent of their value added and then add the same value back to materials.[35] This 17 percent value is estimated by Greenwood et al. (1997) as the equipment share in output. We then pro-rate the equipment share of value added to different materials in proportion to their use within each sector.

We interpret the remaining part of the gross operational surplus in a given sector as compensation for services provided by real estate. We adopt the convention that all land and structures are managed by firms in the real estate sector, which then sell their services to other sectors. Accordingly, for all sectors other than real estate, we reassign the gross operating surplus remaining, after deducting equipment investment, to purchases from the real estate sector. These surpluses are in turn added to the gross operating surplus of real estate.

It follows that, for all sectors j other than real estate,

$$\sum_n r_n H_n^j = 0.$$

and in each of those sectors,

$$\begin{aligned} P_n^{\text{real estate}} M_n^{\text{real estate},j} &= \text{Purchases from real estate by } j \\ &\quad + \text{Operational Surplus of } j \\ &\quad - \text{Equipment Investment by } j. \end{aligned}$$

In contrast, in the real estate sector,

$$\begin{aligned} \sum_n r_n H_n^{\text{real estate}} &= \text{Total Operational Surplus across all } j \\ &\quad - \text{Total Equipment Investment across all } j. \end{aligned}$$

One can verify that those reassignments do not affect aggregate operational surplus (net of equipment investment), aggregate labor compensation, and aggregate value added (net of equipment investment).

We assume that tradable sectors have a $\gamma_n^j = \gamma^j$, constant across cities and similarly for $\gamma_n^{jj'}$'s. The two non-tradable sectors have city specific parameters. Given these adjustments to the Use Table, the share parameters for tradable sectors follow immediately,

$$\gamma^j = \frac{\sum_{n,k} w_n^k L_n^{kj} + \sum_n r_n H_n^j}{\sum_n (1 + s_n^j) X_n^j}, \ \gamma^{j'j} = \frac{\sum_n P_n^{j'} M_n^{j'j}}{\sum_n (1 + s_n^j) X_n^j}. \tag{41}$$

[35] When gross operation surplus amounts to less than 17 percent of value added, the entire operational surplus is deducted.

Furthermore, we have that, for the non-tradable sectors,

$$\gamma_n^j = \frac{\sum_k w_n^k L_n^{kj} + r_n H_n^j}{(1+s_n^j) X_n^j}, \tag{42}$$

where s_n^j is an ad-valorem subsidy for city n, sector j, which we introduce to account for the fact that part of the sectoral value added calculated by the BEA is in fact paid out in indirect taxes. Finally, since we do not observe use of materials by individual sectors in each city, we assume that, the proportions of materials used in each city by nontradable sectors if fixed at the national averages $\frac{\gamma_n^{j'j}}{1-\gamma_n^j}$ is the same for all cities and satisfies equals $\frac{\sum_n P_n^{j'} M_n^{j'j}}{\sum_{n,j'} P_n^{j'} M_n^{j'j}}$

$$\frac{\gamma_n^{j'j}}{1-\gamma_n^j} = \hat{\gamma}^{j'j} = \frac{\sum_n P_n^{j'} M_n^{j'j}}{\sum_{n,j'} P_n^{j'} M_n^{j'j}}$$

The calibration of γ_n^j and, therefore of $\gamma_n^{j'j}$, will require choosing additional parameters as described below but consistent with the above equations.

Trade costs We assign trade costs to be a log-linear function of distance, that is,

$$\kappa_{nn'}^j = (d_{nn'})^{t^j}$$

where $\kappa_{n,n'}$ is the amount of goods that need to be shipped from location n' in order for one unit of the good to be available in location n, and $d_{n,n'}$ is the distance (in miles) between the two locations.

From equation (36) we can write

$$\log(\pi_{nn'}^j) = -\theta t^j \log(d_{nn'}) + c_n + c_{n'} \tag{43}$$

where c_n and $c_{n'}$ are n and n' location-specific factors

$$c_{n'} = -\theta \log(x_{n'}^j)$$

and

$$c_n = -\log\left(\sum_{n''} \left[\kappa_{nn''}^j x_{n''}^j\right]^{-\theta}\right)$$

We assign trade costs to industries using three different methods. First, we assign two industries (retail, construction and utilities, and real estate) to be non-tradable. Second, we use the estimates in Table 1 of Anderson et al. (2014) to obtain gravity coefficients for services. Third, we rely on equation (43) to obtain the gravity coefficients. In order to do this, we use the 2012 Commodity Flow Survey Public Use Microdata File. We add up shipment values by industry, origin and destination and then, for each industry, we regress the log of those averages on log of average shipment distance between each origin and destination.

The gravity coefficients used are summarized in Table 9, below.

Table 9: Estimated Gravity coefficients ($-\theta t^j$)

Industry	Gravity Coefficient	Source
Retail, Construction and Utilities	$-\infty$	Non-tradable
Food and Beverage	-1.24	own estimate
Textiles	-0.88	own estimate
Wood, Paper, and Printing	-1.36	own estimate
Oil, Chemicals, and Nonmetalic Minerals	-1.32	own estimate
Metals	-1.20	own estimate
Machinery	-0.81	own estimate
Computer and Electric	-0.77	own estimate
Electrical Equipment	-0.64	own estimate
Motor Vehicles (Air, Cars, and Rail)	-0.90	own estimate
Furniture and Fixtures	-1.18	own estimate
Miscellaneous Manufacturing	-0.83	own estimate
Wholesale Trade	-0.563	Anderson et al. (2014)
Transportation and Storage	-0.617	Anderson et al. (2014)
Professional and Business Services	-0.928	Anderson et al. (2014)
Other	-0.724	Anderson et al. (2014)
Communication	-0.297	Anderson et al. (2014)
Finance and Insurance	-0.678	Anderson et al. (2014)
Real Estate	$-\infty$	Non-tradable
Education	-1.01	Anderson et al. (2014)
Health	-1.42	Anderson et al. (2014)
Accomodation	-0.927	Anderson et al. (2014)

See text for own estimation details. Coefficients from Anderson et al. (2014) are extracted from Table 1.

B.2 Model Inversion for the Granular Parameters

From the ACS, we obtain data pertaining to w_n^k, and $\frac{L_n^{kj}}{\sum_{k'} L_n^{k'j}}$. The spatial distribution of CNR shares (L_n^k/L_n) is depicted in Figure 20 below. The Census County Business Patterns (CBP) provide us measures of total employment $\sum_{k'} L_n^{k'j}$ that better match BEA industry-level counts, which we combine with the ACS data to obtain L_n^{kj}. From the BEA, we obtain regional price parity (RPP) indices for each city, disaggregated into goods, services and rents. As explained below, we use the level of rents and the relative price of goods and services, providing us with $2(N-1)$ additional restrictions (we deduct 2 since prices in any given city are only defined relative to those in other cities). Furthermore, we can use the BEA Use Tables to calculate the national share of income from land and structures in the production of real estate, providing us with one additional equation. Lastly, as we explain in more detail below, we can apply J normalizations to each sector.

The data plus normalizations above impose $NK + NKJ + 2N + J - 1$ independent restrictions that allows us to solve for NK values for amenity parameters, A_n^k, NKJ scaling factors in production, $T_n^{kj} \equiv \left(H_n^j\right)^{\gamma_n^j \beta_n^j} \left(\lambda_n^{kj}\right)^{\gamma_n^j (1-\beta_n^j)}$, $(N-1)$ shares of non-residential structures in value added in the real estate sector, $\beta_n^{\text{real estate}}$, $(N-1)$ shares of value added in non-tradable output, and $J-1$ independent values for consumption share parameters, α^j.[36]

[36] Furthermore, additional restrictions imposed on α^j and β^j, specifically that $\alpha^j \in [0,1]$, and $\beta^j \in [0,1]$, imply some overidentifying restrictions.

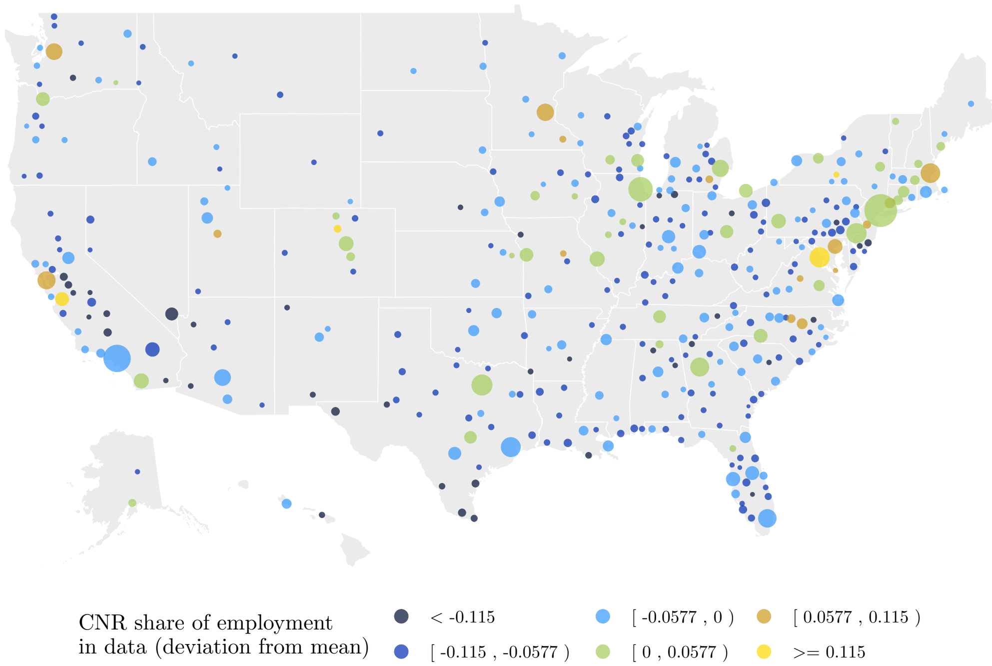

Figure 20: CNR share from 2011-2015

Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city.

The steps below describe the model inversion.

1. *Computing consumption shares*, α^j. We first add up equation (26) across n and j. We then use the factor demand equations (31) and (32) to obtain $\gamma_n^j X_n^j = \left(\sum_n r_n H_n^j + \sum_{n,k} w_n^k L_n^{kj}\right)$, and the national accounting identity (37) to substitute out X_n^j's and C_n^k's from the aggregated equation (26):

$$\alpha^j \sum_{j'} \left(\sum_n r_n H_n^{j'} + \sum_{n,k} w_n^k L_n^{kj'} \right) + \sum_{n,j'} P_n^j M_n^{jj'} = \sum_n r_n H_n^j + \sum_{n,k} w_n^k L_n^{kj} + \sum_{n,j'} P_n^{j'} M_n^{j'j}.$$

The ACS does not provide data on $r_n H_n^j$, $\sum_{j'} P_n^j M_n^{jj'}$ or their sum across cities. While the BEA provides data on sectoral aggregates, those cover the whole country as opposed to only MSA's. Thus, we rely instead on ratios, $\frac{\sum_n r_n H_n^j}{\sum_{n,k} w_n^k L_n^{kj}}$ and $\frac{\sum_{n,j'} P_n^{j'} M_n^{jj'}}{\sum_{n,k} w_n^k L_n^{kj}}$ obtained from the Use Tables, which we can then combine with data on $\sum_{n,k} w_n^k L_n^{kj}$ and the above equation. Specifically,

$$\begin{aligned} &\alpha^j \sum_{n,j',k} w_n^k L_n^{kj'} \left(\frac{\sum_n r_n H_n^{j'}}{\sum_{n,k} w_n^k L_n^{kj'}} + 1 \right) + \sum_{j',n,k} \frac{\sum_{n,j'} P_n^{j'} M_n^{jj'}}{\sum_{n,k} w_n^k L_n^{kj}} w_n^k L_n^{kj'} \\ &= \left(\sum_{n,k} w_n^k L_n^{kj} \right) \times \left(\frac{\sum_n r_n H_n^j}{\sum_{n,k} w_n^k L_n^{kj}} + \frac{\sum_{n,j'} P_n^{j'} M_n^{j'j}}{\sum_{n,k} w_n^k L_n^{kj}} + 1 \right), \end{aligned}$$

or

$$\alpha^j \sum_{n,j',k} w_n^k L_n^{kj'} \left(\varrho_H^{j'} + 1 \right) + \sum_{j',n,k} w_n^k L_n^{kj'} \varrho_M^{jj'}$$
$$= \sum_{n,k} w_n^k L_n^{kj} \left(\varrho_H^j + \varrho_M^{j'j} + 1 \right),$$

where ϱ_H^j and $\varrho_M^{j'j}$ denote, respectively, the ratio of national aggregate rental income and the ratio of national aggregate material inputs usage from sector j' to national aggregate wage income in sector j which are consistent with the Use Tables. The J equations above can be solved for J values of α^j . One can verify that any value of α^j obtained from those equations will satisfy $\sum_j \alpha^j = 1$. One complicating factor is that in each sector j, α^j must live in $[0, 1]$. However, because of measurement inconsistencies between ACS and BEA data, the procedure generates negative values of α^j in three out of 22 sectors. One of those sectors ("Oil, Chemicals, and Nonmetallic Minerals") indeed has much of its employment located outside of urban areas. We use information from the Use Tables to calibrate α^j in that sector, setting it equal to 5.57 percent. The other two sectors ("Wood, Paper, and Printing", and "Metals") are to a large degree producers of inputs for other industries, so that we set their consumption shares to 0. To ensure that all equations hold while satisfying those restrictions, we allow $\varrho_M^{j'j}$'s to deviate somewhat from those obtained from the Use Tables. This requires adjusting γ^j and $\gamma^{j'j}$ for the tradable sectors, since those satisfy $\gamma^j = \frac{1+\varrho_H^j}{1+\varrho_H^j+\sum^{j'} \varrho_M^{j'j}}$, $\gamma^{j'j} = \frac{\varrho_M^{j'j}}{1+\varrho_H^j+\sum^{j'} \varrho_M^{j'j}}$.

2. *Expressing gross output and rental income from each sector and city as functions of share parameters and wage bills.* Using the labor demand equations (31), we obtain

$$\sum_{n'} \pi_{n'n}^j X_{n'}^j = \frac{\sum_k w_n^k L_n^{kj}}{\left(1 - \beta_n^j\right) \gamma_n^j}. \tag{44}$$

In the non-tradable sectors, $\pi_{nn}^j = 1$ and $\pi_{n'n}^j = 0$ for $n' \neq n$ so that

$$X_n^j = \frac{\sum_k w_n^k L_n^{kj}}{\left(1 - \beta_n^j\right) \gamma_n^j}.$$

For all sectors other than real estate, we have that $\beta_n^j = 0$, so that $r_n H_n^j = 0$. For the real estate sector, we have from the first-order conditions in that sector that

$$r_n H_n^{\text{real estate}} = \frac{\beta_n^{\text{real estate}}}{1 - \beta_n^{\text{real estate}}} \sum_k w_n^k L_n^{k,\text{real estate}}.$$

Since real estate services are the only sector with positive rental income, this is also equal to the total rental income in each city.

3. *Computing the shares of land and structures in value added for the real estate sector, $\beta_n^{\text{real estate}}$.* We use equations (30) to substitute out $P_n C_n^k$ in equations (26). We then apply the relationships from equation (44) to substitute out gross output in (31) to (33), and use the resulting equations to substitute out factor demands in (26). Given that in the non-tradable sectors

("real estate" and "retail, construction, and utilities"), expenditures are equal to gross output, this implies that, for $j \in \{$"real estate","retail, construction, and utilities"$\}$, we have that

$$\frac{1}{1-\beta_n^j}\frac{\sum_k w_n^k L_n^{kj}}{\gamma_n^j}$$
$$= \alpha^j \sum_k w_n^k L_n^k$$
$$+ \alpha^j \sum_k \frac{L_n^k}{L^k} b^k \sum_{n'} \left(\frac{\beta_{n'}^{\text{real estate}}}{1-\beta_{n'}^{\text{real estate}}} \sum_{k'} w_{n'}^{k'} L_{n'}^{k',\text{real estate}} \right)$$
$$+ \sum_{j'} \frac{1-\gamma_n^{j'}}{\gamma_n^{j'}(1-\beta_n^{j'})} \hat{\gamma}^{jj'} \sum_k w_n^k L_n^{kj'},$$

where we are using the fact that $\beta_n^j = 0$ for all sectors other than real estate. Given that we have two non-tradable sectors, this is a system of $2N$ equations, in N values for γ_n^j and N values of $\beta_n^{\text{real estate}}$.

4. *Computing individual values for nominal expenditures, X_n^j, in tradable sectors.* We use equations (30) to substitute out $P_n C_n^k$ in equations (26). We then apply the relationships from equation (44) to substitute out gross output in (31) to (33), and use the resulting equations to substitute out factor demands in (26). In the tradable sectors, this gives us

$$X_n^j = \alpha^j \sum_k w_n^k L_n^k$$
$$+ \alpha^j \sum_k \frac{L_n^k}{L^k} b^k \sum_{n'} \left(\frac{\beta_{n'}^{\text{real estate}}}{1-\beta_{n'}^{\text{real estate}}} \sum_{k'} w_{n'}^{k'} L_{n'}^{k'j'} \right)$$
$$+ \sum_{j'} \gamma_n^{jj'} \frac{\sum_k w_n^k L_n^{kj'}}{\left(1-\beta_n^{j'}\right)\gamma_n^{j'}}$$

Given values for β_n^j from the previous step, values for X_n^j are then immediately determined from the data.

5. *Computing relative cost indices for tradable goods, $\tilde{x}_n^j$.* For $N(J-2)$ tradable sectors (all but "real estate," as well as "retail, construction, and utilities"), we now solve for $(N-1)(J-2)$ values of the cost index, $\frac{x_n^j}{\sum_{n'} x_{n'}^j}$, for each $j \in \{1, ..., J\}$ from the system of $(N-1)(J-2)$ independent equations,

$$\frac{\sum_k w_n^k L_n^{kj}}{\left(1-\beta_n^j\right)\gamma_n^j} = \sum_{n'=1}^{N} \pi_{n'n}^j \left(\mathbf{x}^j\right) X_{n'}^j,$$

where $\mathbf{x}^j = \{x_1^j, ..., x_N^j\}$ is the vector of unit production costs. This system comprises only $(N-1)(J-2)$ independent equations since, for each j, adding up the right-hand-side and left-hand-side over n gives the same result on both sides irrespective of $\mathbf{x}^j$. At the same time $\pi^j_{n'n}\left(\mathbf{x}^j\right)$ is homogeneous of degree 0 in $\mathbf{x}^j$ for each j in equation (36), so that we can still solve for the ratio, $\tilde{x}^j_n \equiv \frac{x^j_n}{\sum_{n'} x^j_{n'}}$. [37]

6. *Computing relative tradable consumer prices, $\widetilde{P}^j_n$, in every sector and city.* Substituting $\tilde{x}^j_n$ from the previous step into equation(34) and rearranging, we have that for the tradable sectors,

$$P^j_n = \Gamma\left(\xi\right)^{\frac{1}{1-\eta}} \left(\sum_{n'} \left[\kappa^j_{nn'}\tilde{x}^j_n\right]^{-\theta}\right)^{-\frac{1}{\theta}} \times \sum_{n'} x^j_{n'},$$

which gives a system of $N(J-2)$ equations. We can thus determine

$$\Xi^j_P \equiv \Gamma\left(\xi\right)^{\frac{1}{1-\eta}} \frac{\sum_{n'} x^j_{n'}}{P^j} = \left[\sum \varpi^j_n \left(\sum_{n'} \left[\kappa^j_{nn'}\tilde{x}^j_n\right]^{-\theta}\right)^{-\frac{1}{\theta}}\right]^{-1}$$

for each j by imposing $\sum_n \varpi^j_n P^j_n = P^j$, where ϖ^j_n are model-consistent expenditure weights given by $X^j_n / \sum_{n'} X^j_{n'}$ obtained in step 4. We may then then obtain for all tradable j's

$$\tilde{P}^j_n \equiv \frac{P^j_n}{P^j} = \Xi^j_P \left(\sum_{n'} \left[\kappa^j_{nn'}\tilde{x}^j_n\right]^{-\theta}\right)^{-\frac{1}{\theta}}.$$

Note that data on P^j is only available in changes from a base period. Thus, we define the base period to be 2011-2015, our benchmark period, and set $P^j = 1$ in all sectors in that period. For the remainder of the analysis, therefore, $\tilde{P}^j_n = P^j_n$.

7. *Computing non-tradable consumer prices.* In the non-tradable sectors, we have that $P^j_n = \Gamma\left(\xi\right)^{\frac{1}{1-\eta}} x^j_n$ for all n and j, and for those sectors, we determine prices based on data from the Regional Price Parity (RPP) indices calculated by the BEA. We directly obtain values for $\Gamma\left(\xi\right)^{\frac{1}{1-\eta}} x^{\text{real estate}}_n \equiv P^{\text{real estate}}_n$ from the RPP estimates of the price of real estate services in different cities. For the other non-tradables ("retail, construction, and utilities"), we choose $P^{\text{retail, etc.}}_n = \Gamma\left(\xi\right)^{\frac{1}{1-\eta}} x^{\text{retail, etc.}}_n$ so that the price of services (other than real estate) relative to tradable goods in the model matches its counterpart in the RPP. To carry out this calculation, observe that the price index for services can be defined by $P^{\text{Services}}_n = \Pi_{j\in\text{Services}} \left(\frac{\Sigma_{j'\in\text{Services}}\alpha^{j'} P^j_n}{\alpha_j}\right)^{\frac{\alpha^j}{\Sigma_{j'\in\text{Services}}\alpha^{j'}}}$ where the service sectors include retail, etc., wholesale trade, transportation and storage, professional and business services, other, communication, finance and insurance, education, health, and accommodation. The price index for goods can be defined analogously where the goods sector includes all remaining sectors other than real estate.

8. *Computing firm productivity in different sectors, j, and cities, n, associated with occupation k, λ^{kj}_n.* From equations (31), we have that

[37]Numerically, the system is easier to solve for $\frac{\left(x^j_n\right)^{\theta j}}{\sum_{n'}\left(x^j_{n'}\right)^{\theta j}}$ from which we can easily obtain values for x^j_n's.

$$w_n^k L_n^{kj} = \frac{\left(\frac{w_n^k}{\lambda_n^{kj}}\right)^{1-\epsilon}}{\sum_{k'}\left(\frac{w_n^{k'}}{\lambda_n^{k'j}}\right)^{1-\epsilon}} \sum_k w_n^k L_n^{kj},$$

which can rewrite as

$$w_n^k L_n^{kj} = \frac{\left(\frac{\sum_{k'}\lambda_n^{k'j}}{\lambda_n^{kj}} w_n^k\right)^{1-\epsilon}}{\sum_{k'}\left(\frac{\sum_{k'}\lambda_n^{k'j}}{\lambda_n^{k'j}} w_n^{k'}\right)^{1-\epsilon}} \sum_k w_n^k L_n^{kj}.$$

Thus, for each city n and industry j, we can use $K-1$ of those equations to solve for $K-1$ ratios, $\tilde{\lambda}_n^{kj} = \frac{\lambda_n^{kj}}{\sum_{k'}\lambda_n^{k'j}}$. With some rearrangement, those can be written as

$$\tilde{\lambda}_n^{kj} = \frac{\left(w_n^k\right)^{\frac{\epsilon}{\epsilon-1}}\left(L_n^{kj}\right)^{\frac{1}{\epsilon-1}}}{\sum_{k'}\left(w_n^{k'}\right)^{\frac{\epsilon}{\epsilon-1}}\left(L_n^{k'j}\right)^{\frac{1}{\epsilon-1}}}.$$

From equations (45) (obtained by substituting the factor demand equations (31), (32) and (33) into equations (35)), and the value for $r_n H_n$ obtained in step 2, we obtain

$$\begin{aligned} & H_n^{\gamma_n^j\beta_n^j}\left(\sum_{k'}\lambda_n^{k'j}\right)^{\gamma_n^j(1-\beta_n^j)}\Gamma(\xi)^{-\frac{1}{1-\eta}} \\ &= \frac{B_n^j}{\tilde{x}_n^j \Xi_P^j}\left\{\left(\sum_j \frac{\beta_n^j}{1-\beta_n^j}\sum_k w_n^k L_n^{kj}\right)^{\beta_n^j}\left[\sum_{k=1}^K \left(\frac{w_n^k}{\tilde{\lambda}_n^{kj}}\right)^{1-\epsilon}\right]^{\frac{1-\beta_n^j}{1-\epsilon}}\right\}^{\gamma_n^j}\prod_{j'=1}^J\left(P_n^{j'}\right)^{\gamma_n^{j'j}}, \end{aligned} \tag{45}$$

where we set $\Gamma(\xi)^{-\frac{1}{1-\eta}} = 1$ since it is common to all sectors and cities and thus immaterial in any counterfactual exercise. Recall that the use of land and structures as inputs has been folded in the real estate sector that then sells real estate services to all other sectors (i.e. $\beta_n^j = 0$ in all sectors but real estate). Then, multiplying both sides of equation (45) by the ratios $\left(\tilde{\lambda}_n^{kj}\right)^{\gamma_n^j}$ gives $NK(J-1)$ values for the productivity of firms in different sectors, j, and cities, n, associated with occupation k, $T_n^{kj} \equiv \left(H_n^j\right)^{\gamma_n^j\beta_n^j}\left(\lambda_n^{kj}\right)^{\gamma_n^j(1-\beta_n^j)}$, which, in the special case where one abstracts from differences in occupational composition across cities, reproduces measured regional and sectoral productivity in Caliendo et al. (2018).

9. *Computing the idiosyncratic amenity distribution parameter ν and amenity shifters A_n^k for each city n and occupation k*

   To compute ν we match the estimate for local labor supply elasticity with respect to local wage estimated by Fajgelbaum, Serrato and Zidar (2018) of 1.14. In our setup, for any occupation k and city n, this elasticity is $\nu\frac{w_n^k}{P_n C_n^k}$. The average elasticity is equal to 1.14 if $\nu = 2.017$. Given ν, we can now back-out amenities from the labor supply equation 2.[38]

[38]For given k, labor supply is homogeneous of degree zero in A_n^k. This implies that amenities are only determined

B.3 Instrumenting for Employment Level and Composition

In order to isolate the residual simultaneity between exogenous productivity variation and labor allocation, we resort to variants of instruments proposed in the literature. Specifically, we follow Ciccone and Hall (1996) and use population a century prior to our data period to capture historical determinants of current population, and we follow Card (2001) and Moretti (2004a), and use variation in early immigrant population and the presence of land-grant colleges to capture historical determinants of skill composition of cities. We now discuss the particular instruments in more detail.

Population in 1920 Ciccone and Hall (1996) argue for the validity of historical variables as instruments under the assumption that, after allowing for the controls described above, original sources of agglomeration only affect current population patterns through the preferences of workers, and not through their effect on the residual component of productivity. This reasoning motivates using population almost one hundred years prior to our data period as an instrument and will also serve as motivation for the other instruments, described below.

Irish immigration in 1920 Next, we use the fraction of Irish immigrants in the population of each city in 1920. This instrument is motivated by Card (2001), who uses the location of immigrant communities as an instrument for labor supply in different occupations. For our purposes, we focus on the location of Irish immigrants following evidence reviewed by Neal (1997), and further studied by Altonji et al. (2005), showing that attending catholic schools substantially increases the likelihood of completing high school and college education. We use as an instrument the fraction of Irish immigrants, rather than the overall catholic population, because Irish immigrants represented the first wave of catholic immigration to the U.S. and, therefore, historically were the first to invest in education. As additional validation for this instrument, we compile data on the current location of catholic colleges, and observe that MSAs in which catholic colleges are present had in 1920 more than three times the fraction of Irish immigrants as other locations.

The Presence of land-grant colleges Lastly, following Moretti (2004a), we also use as an instrument the presence of a land-grant college within the city. Land-grant colleges were established as a result of the Morrill Act of 1862, and extended in 1890, a federal act that sought to give states the opportunity to establish colleges in engineering and other sciences. Since the act is more than a century old, the presence of a land-grant college in the city is unlikely to be related to unobservable factors affecting productivity in different cities over our base period, $2011-2015$. At the same time, as shown in Moretti (2004a), the presence of land-grant colleges is generally correlated with the composition of skills across cities.

B.4 A check on the instruments: Estimates in Counterfactual without Externalities

Table 10 below shows the results from carrying out the same estimation exercises as in Table 7 using employment and productivity values obtained from a counterfactual equilibrium in which externality elasticities, $\tau^{R,k}$ and $\tau^{L,k}$, are set to zero but all other model parameters are kept at their original levels.

up to an arbitrary, occupation-specific scaling constant, that is, we can change $A_n|^k$ to $\tilde{A}_n^k$ so that $\tilde{A}_n^k = m^k A_n^k$ without any observable implications.

Table 10: Estimates with data generated by counterfactual without externalities

	(1) **OLS**		(2) **2SLS**		(3) **CUE**	
VARIABLES	CNR	non-CNR	CNR	non-CNR	CNR	non-CNR
$\gamma_n^j log(\frac{L_n^k}{L_n})$	-0.986***	-0.167	-0.276	-0.917	-0.0660	0.0778
	(0.29)	(0.47)	(0.72)	(0.94)	(0.72)	(0.93)
$\gamma_n^j log(L_n)$	0.0174	-0.0371	-0.0267	-0.0693	0.00385	0.00174
	(0.06)	(0.04)	(0.06)	(0.04)	(0.06)	(0.04)
Observations	7,460	7,460	7,460	7,460	7,460	7,460
K.P. F			7.445	7.516	7.445	7.516
S.W.F. L_n^k Share			11.77	11.83	11.77	11.83
S.W.F. L_n			14.64	16.96	14.64	16.96

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

B.5 Model-Implied IV

In this exercise, we estimate externalities using an IV implied by the model. This is obtained by calculating the counter-factual allocation associated with an economy where, for any given industry/occupation category, productivity is constant across cities, and using the resulting counter-factual labor allocation as instruments.

This instrument will correct for a reverse causality problem since, by construction, there is no exogenous variation in productivity across cities. Table 11 below shows the estimates. The F-statistics are very large, implying no need to explore GMM-CUE estimates. Moreover, the coefficients present the same general pattern as our baseline estimates: occupational externalities are generally stronger than those associated with total population.

C The Planner's Problem

This section describes the solution to the planner's problem taking as given that workers in different occupations can freely choose in which city to live. Under this assumption, the expected utility of a worker of type k is given by equation (16). Given welfare weights for each occupation, ϕ^k, the utilitarian planner then solves

$$\mathcal{W} = \sum_k \phi^k U \left[\Gamma \left(\frac{\nu - 1}{\nu} \right) \left(\sum_{n=1}^{N} \left(A_n^k C_n^k \right)^{\nu} \right)^{\frac{1}{\nu}} \right] L^k, \tag{46}$$

where recall that C_n^k aggregates final goods from different sectors:

$$C_n^k = \prod_j \left(C_n^{kj} \right)^{\alpha^j}. \tag{47}$$

The planner maximizes (46) subject to the resource constraints for final goods,

Table 11: Externality estimates with Model Implied IV's

	(1)		(2)	
	OLS		**2SLS**	
VARIABLES	CNR	non-CNR	CNR	non-CNR
$\gamma_n^j log(\frac{L_n^k}{L_n})$	0.889***	0.702***	0.752***	0.454**
	(0.12)	(0.22)	(0.12)	(0.21)
$\gamma_n^j log(L_n)$	0.386***	0.322***	0.405***	0.308***
	(0.05)	(0.04)	(0.05)	(0.04)
Observations	7,460	7,460	7,460	7,460
K.P. F			600.2	595.5
S.W.F. L_n^k Share			1187	1205
S.W.F. L_n			1705	1778

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

$$\sum_k L_n^k C_n^{kj} + \sum_{j'} \int M_n^{jj'}(\mathbf{z}) d\Phi(\mathbf{z}) = \left(\int \left[\sum_{n'} Q_{nn'}^j(\mathbf{z}) \right]^{\frac{\eta-1}{\eta}} d\Phi(\mathbf{z}) \right)^{\frac{\eta}{\eta-1}}, \tag{48}$$

where $Q_{nn'}^j(\mathbf{z})$ are the purchases of intermediate goods produced in city n' by final goods firms in city n, the resource constraints for intermediate goods of all varieties $\mathbf{z}$ and industries j produced in all cities n

$$\sum_{n'} Q_{n'n}^j(\mathbf{z}) \kappa_{n'n}^j = q_n^j(\mathbf{z}), \quad \forall \mathbf{z} \in \mathbb{R}_n^+, \tag{49}$$

where

$$q_n^j(\mathbf{z}) = z_n \left[H_n^j(\mathbf{z})^{\beta_n^j} \left[\sum_k \left(\lambda_n^{kj}(\mathbf{L}_n) L_n^{kj}(\mathbf{z}) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}(1-\beta_n^j)} \right]^{\gamma_n^j} \prod_{j'} M_n^{j'j}(\mathbf{z})^{\gamma_n^{j'j}},$$

labor markets constraints in all locations,

$$\sum_j \int L_n^{kj}(\mathbf{z}) d\Phi(\mathbf{z}) = L_n^k, \tag{50}$$

where labor supply in each city, L_n^k, satisfies

$$L_n^k = \frac{\left(A_n^k C_n^k\right)^\nu}{\sum_{n'} \left(A_{n'}^k C_{n'}^k\right)^\nu} L^k, \tag{51}$$

the resource constraints in the use of land and structures,

$$\sum_j \int H_n^j(\mathbf{z}) d\Phi(\mathbf{z}) = H_n, \tag{52}$$

as well as non-negativity constraints applying to both household consumption of different goods and input flows:

$$C_n^{kj} \geq 0 \text{ and } Q_{n'n}^j(\mathbf{z}) \geq 0.$$

From the resource constraint on local labor markets (50), and the labor supply condition (51), it follows immediately that national labor markets clear (i.e., $\sum_{n,j} \int L_n^{kj}(\mathbf{z}) d\Phi(\mathbf{z}) = L^k$).

C.1 Solving the Planner's Problem

We solve the Planner's problem for interior allocations, (i.e., where C_n^k and L_n^k are strictly greater than zero for all n and k). For each city n and sector j, let P_n^j be the Lagrange multiplier corresponding to the final goods resource constraint in city n, sector j (48), $\widetilde{P}_n$ the multiplier corresponding to the aggregation of sectoral goods in each city (47), and $\widetilde{p}_n^j(\mathbf{z})$ the multiplier corresponding to the intermediate goods resource constraints (49). For each city n and occupation k, let w_n^k be the multiplier corresponding to regional labor market clearing (50), W_n^k the multiplier corresponding to the definitions of employment in each occupation and sector (51). Finally, for each city n, let r_n denote the multiplier corresponding to market clearing for structures (52).

The first-order conditions associated with the planner's problem are:

$$\partial C_n^{kj} : \ \widetilde{P}_n \alpha^j \frac{C_n^k}{C_n^{kj}} = P_n^j L_n^k, \tag{53}$$

which also defines an ideal price index,

$$P_n = \frac{\widetilde{P}_n}{L_n^k} = \prod_j \left(\frac{P_n^j}{\alpha^j} \right)^{\alpha^j}. \tag{54}$$

In addition,

$$\begin{aligned} \partial C_n^k \quad &: \quad \phi^k U'\left(v^k\right) v^k \frac{\left(A_n^k C_n^k\right)^\nu}{\sum_{n'} \left(A_{n'}^k C_{n'}^k\right)^\nu} \frac{1}{C_n^k} L^k \\ &= \quad L_n^k P_n - \sum_{n'=1}^N \frac{\partial \zeta_{n'}^k\left(\mathbf{C}^k\right)}{\partial C_n^k} W_{n'}^k, \end{aligned} \tag{55}$$

where

$$v^k = \Gamma\left(\frac{\nu - 1}{\nu}\right) \left(\sum_{n'} \left(A_{n'}^k C_{n'}^k\right)^\nu \right)^{\frac{1}{\nu}}$$

and

$$\frac{\partial \zeta_{n'}^k\left(\mathbf{C}^k\right)}{\partial C_n^k} = \begin{cases} \left(\frac{\nu}{C_n^k}\right)\left(1 - \frac{L_n^k}{L^k}\right) L_n^k & \text{if } n' = n \\ -\left(\frac{\nu}{C_n^k}\right)\left(\frac{L_{n'}^k}{L^k}\right) L_n^k & \text{if } n' \neq n \end{cases}.$$

Also

$$\partial L_n^k : \sum_{j=1}^{J} P_n^j C_n^{kj} - \widetilde{w}_n^k + W_n^k = 0. \tag{56}$$

where

$$\begin{aligned}\widetilde{w}_n^k \;&=\; w_n^k \\ &+\sum_j \int \frac{\partial z_n \left[H_n^j(\mathbf{z})^{\beta_n^j} \left[\sum_{k''} \left(\lambda_n^{k''j}\left(\mathbf{L}_n\right) L_n^{k''j}\left(\mathbf{z}\right) \right)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}\left(1-\beta_n^j\right)} \right]^{\gamma_n^j} \prod_{j'=1}^{J} M_n^{j'j}(\mathbf{z})^{\gamma_n^{j'j}}}{\partial L_n^k} \widetilde{p}_n^j\left(\mathbf{z}\right) d\mathbf{z}\end{aligned} \tag{57}$$

denotes the total social marginal value of an extra worker of type k in city n. On the production side, efficient allocations dictate

$$\partial Q_{nn'}^j(\mathbf{z}) : \begin{cases} Q_{nn'}^j(\mathbf{z}) > 0 & \text{if } \kappa_{nn'}^j \widetilde{p}_{n'}^j(\mathbf{z}) = P_n^j \left(Q_n^j\right)^{\frac{1}{\eta}} \left[\sum_{n'=1}^N Q_{nn'}^j(\mathbf{z})\right]^{-\frac{1}{\eta}} d\Phi(\mathbf{z}) \\ Q_{nn'}^j(\mathbf{z}) = 0 & \text{if } \kappa_{nn'}^j \widetilde{p}_{n'}^j(\mathbf{z}) > P_n^j \left(Q_n^j\right)^{\frac{1}{\eta}} \left[\sum_{n'=1}^N Q_{nn'}^j(\mathbf{z})\right]^{-\frac{1}{\eta}} d\Phi(\mathbf{z}) \end{cases}. \tag{58}$$

This last equation delivers efficient trade shares, $\pi_{nn'}^j$, and prices, P_n^j, using the usual Eaton and Kortum derivations. In addition,

$$\partial L_n^{kj}(\mathbf{z}) : \gamma_n^j (1-\beta_n^j) \frac{q_n^j(\mathbf{z})}{L_n^{kj}\left(\mathbf{z}\right)} \frac{\left(\frac{w_n^k}{\left(\lambda_n^{kj}(\mathbf{L}_n)\right)}\right)^{1-\epsilon}}{\sum_{k'} \left(\frac{w_n^{k'}}{\left(\lambda_n^{kj}(\mathbf{L}_n)\right)}\right)^{1-\epsilon}} \widetilde{p}_n^j(\mathbf{z}) = w_n^k d\Phi(\mathbf{z}), \tag{59}$$

$$\partial H_n^j(\mathbf{z}) : \gamma_n^j \beta_n^j \frac{q_n^j(\mathbf{z})}{H_n^j(\mathbf{z})} \widetilde{p}_n^j(\mathbf{z}) = r_n d\Phi(\mathbf{z}), \tag{60}$$

$$\partial M_n^{j'j}(\mathbf{z}) : \gamma_n^{j'j} \frac{q_n^j(\mathbf{z})}{M_n^{j'j}(\mathbf{z})} \widetilde{p}_n^j(\mathbf{z}) = P_n^{j'} d\Phi(\mathbf{z}). \tag{61}$$

With the usual manipulations of these equations, one obtains

$$\widetilde{p}_n^j(\mathbf{z}) \equiv p_n^j(\mathbf{z}) d\Phi(\mathbf{z}) = \frac{x_n^j d\Phi(\mathbf{z})}{z_n}, \tag{62}$$

where

$$x_n^j = B_n^j \left[r_n^{\beta_n^j} \left[\sum_k \left(\frac{w_n^k}{\left(\lambda_n^{kj}\left(\mathbf{L}_n\right)\right)} \right)^{1-\epsilon} \right]^{\frac{1-\beta_n^j}{1-\epsilon}} \right]^{\gamma_n^j} \prod_{j'} \left(P_n^{j'}\right)^{\gamma_n^{j'j}}, \tag{63}$$

and B_n^j is defined as above.

D Characterization of the Planner's Solution

In the decentralized equilibrium, the budget constraint of a household of type k in city n satisfies

$$P_n C_n^k = w_n^k + \chi^k,$$

where $\chi^k = b^k \frac{\sum_{n'} r_{n'} H_{n'}}{\sum_{n',j} L_{n'}^{k,j}}$. In contrast, we now show that the consumption of a household of type k in city n implied by the planner's solution satisfies

$$P_n C_n^k = \frac{\nu}{1+\nu} \widetilde{w}_n^k + \chi^k + R^k,$$

and recall that $\widetilde{w}_n^k$ is the social marginal product of labor associated with occupation k in city n.

Proof:
Equation (55) may alternatively be expressed as

$$\phi^k U'(v^k) v^k \frac{L_n^k}{C_n^k} = L_n^k P_n - \left(\frac{\nu}{C_n^k}\right) L_n^k W_n^k + \sum_{n'=1}^{N} \left(\frac{\nu}{C_n^k}\right) \left(\frac{L_{n'}^k}{L^k}\right) L_n^k W_{n'}^k,$$

where v^k is defined in equation (16). Alternatively, we have that

$$\underbrace{\phi^k U'(v^k) v^k - \sum_{n'} \nu \left(\frac{L_{n'}^k}{L^k}\right) W_{n'}^k}_{(1+\nu)(\chi^k + R^k)} = P_n C_n^k - \nu W_n^k.$$

Substituting for W_n^k from (56) in this last expression gives

$$P_n C_n^k = \nu \left(\widetilde{w}_n^k - P_n C_n^k\right) + (1+\nu)\left(\chi^k + R^k\right)$$

or

$$P_n C_n^k = \frac{\nu}{1+\nu} \widetilde{w}_n^k + \chi^k + R^k. \tag{64}$$

□

Observe that we can also use (56) to write $\chi^k + R^k$ as a function of prices, $\widetilde{w}_n^k$, P_n, and consumption, C_n^k. In particular,

$$\chi^k + R^k = \frac{\phi^k U'(v^k) v^k}{1+\nu} - \frac{\nu}{1+\nu} \sum_{n'} \left(\frac{L_{n'}^k}{L^k}\right) \left(\widetilde{w}_n^k - P_{n'} C_{n'}^k\right).$$

We can then obtain an expression for the total consumption expenditures of households of type k by adding (64) across cities n, with the expression for $\chi^k + R^k$ substituted in,

$$\phi^k U'(v^k) v^k L^k = \sum_n P_n C_n^k L_n^k. \tag{65}$$

Substituting out $\phi^k U'(v^k) v^k$ back into the expression for $\chi^k + R^k$ and rearranging, we obtain

$$\chi^k + R^k = \frac{\sum_n P_n C_n^k L_n^k}{L^k} - \sum_n \frac{\nu^k}{1+\nu} \left(\frac{L_n^k}{L^k}\right) \widetilde{w}_n^k.$$

Finally, note that $\sum_{n,k} P_n C_n^k L_n^k = \sum_{n,k} \left(w_n^k L_n^k + r_n H_n \right)$, so that

$$\sum_k L^k(\chi^k + R^k) = \sum_{n,k} \frac{1}{1+\nu} w_n^k L_n^k - \sum_{n,k} \frac{\nu}{1+\nu} \left(\widetilde{w}_n^k - w_n^k \right) L_n^k + \sum_n r_n H_n \tag{66}$$

The individual values for χ^k are determined to be such that equation (66) is satisfied.

D.1 The Social and Private Marginal Value of Workers of type k in city n (Proof of Lemma 1)

Solving the derivative in the equation defining the social value of workers of type k in city n (57), we obtain

$$\begin{aligned} &\tilde{w}_n^k - w_n^k \\ = &\sum_j \int \frac{\partial z_n^j \left[H_n^j(\mathbf{z})^{\beta_n^j} \left[\sum_{k'} \left(\lambda_n^{k'j}\left(\mathbf{L}_n\right) L_n^{k'j}\left(\mathbf{z}\right) \right)^{1-\frac{1}{\epsilon}} \right]^{\frac{\epsilon}{1-\epsilon}\left(1-\beta_n^j\right)} \right]^{\gamma_n^j} \prod_{j'=1}^{J} M_n^{j'j}(\mathbf{z})^{\gamma_n^{j'j}}}{\partial L_n^k} p_n^j(\mathbf{z}) d\Phi(\mathbf{z}) \end{aligned}$$

where $p_n^j(\mathbf{z}) d\Phi(\mathbf{z}) = \tilde{p}_n^j(\mathbf{z})$. This expression is equivalent to

$$\tilde{w}_n^k - w_n^k = \sum_{j,k'} (1-\beta_n^j)\gamma_n^j \frac{\left(\frac{w_n^{k'}}{\lambda_n^{k'j}(\mathbf{L}_n)} \right)^{1-\epsilon}}{\sum_{k''} \left(\frac{w_n^{k''}}{\lambda_n^{k'j}(\mathbf{L}_n)} \right)^{1-\epsilon}} \frac{1}{\lambda_n^{k'j}\left(\mathbf{L}_n\right)} \frac{\partial \lambda_n^{k'j}\left(\mathbf{L}_n\right)}{\partial L_n^k} q_n^j(\mathbf{z}) p_n^j(\mathbf{z}) d\Phi(\mathbf{z}).$$

Rearranging and integrating equation (59) yields

$$w_n^k L_n^{kj} = (1-\beta_n^j)\gamma_n^j \frac{\left(\frac{w_n^{k}}{\lambda_n^{kj}(\mathbf{L}_n)} \right)^{1-\epsilon}}{\sum_{k'} \left(\frac{w_n^{k'}}{\lambda_n^{k'j}(\mathbf{L}_n)} \right)^{1-\epsilon}} \int q_n^j(\mathbf{z}) p_n^j(\mathbf{z}) d\Phi(\mathbf{z}),$$

so that the expression for the deviation of private from social marginal product of labor simplifies further to

$$\tilde{w}_n^k - w_n^k = \sum_{j,k'} w_n^{k'} \frac{L_n^{k'j}}{L_n^k} \frac{\partial \ln \lambda_n^{k'j}\left(\mathbf{L}_n\right)}{\partial \ln L_n^k}. \tag{67}$$

D.2 Implementation (Proof of Proposition 1)

We now discuss the implementation of the optimal policy. One possible implementation is to combine a direct employment subsidy to firms that is specific to cities and occupations (Δ_n^k), a linear occupation-specific labor income tax (t_L^k), combined with occupation-specific transfers (R^k).

With externalities in occupations, the social and private marginal products of labor differ. The first step in the implementation of optimal allocations, therefore, is to subsidize firms in different locations to hire different occupation types. We define $\widetilde{w}_n^k$ to be the after-subsidy wage associated with workers in occupation k living in city n such that

$$\widetilde{w}_n^k = w_n^k + \Delta_n^k,$$

where Δ_n^k is a per-worker subsidy offered to firms in city n hiring workers of type k. With these subsidized wages in place, we take advantage of various additional taxes and transfers to implement optimal allocations. In particular, equation (1) becomes

$$I_n^k = (1 - t_L^k)\widetilde{w}_n^k + \chi^k + R^k, \; (NK \text{ eqs.}) \,. \tag{68}$$

where transfers have to be such that the government budget balances,

$$\sum_{n,k} L_n^k R^k = \sum_{n,k} t_L^k w_n^k L_n^k - \sum_{n,k} (1 - t_L^k)\Delta_n^k L_n^k. \tag{69}$$

We also have that labor demand depends only on pre-subsidy wages, w_n^k,

$$w_n^k L_n^{kj}(\mathbf{z}) = \frac{\left(\frac{w_n^k}{\lambda_n^{kj}(\mathbf{L_n})}\right)^{1-\epsilon}}{\sum_{k'}\left(\frac{w_n^k}{\lambda_n^{k'j}(\mathbf{L_n})}\right)^{1-\epsilon}} \gamma_n^j \left(1 - \beta_n^j\right) p_n^j(\mathbf{z}) q_n^j(\mathbf{z}), \; (NKJ \; eqs.) \tag{70}$$

$$x_n^j = B^j \left[r_n^{\beta_n^j} \left[\sum_k \left(\frac{w_n^k}{\lambda_n^{kj}\left(\mathbf{L}_n\right)} \right)^{1-\epsilon} \right]^{\frac{1-\beta_n^j}{1-\epsilon}} \right]^{\gamma_n^j} \prod_{j'} \left(P_n^{j'}\right)^{\gamma_n^{j'j}}, \tag{71}$$

Definition 1. An equilibrium with taxes and transfers is defined as the equilibrium without taxes and transfers but with the additional conditions that i) I_n^k is given by equation (68), ii) the first-order condition describing intermediate goods producers' labor demand is given by (70), iii) the cost index x_n^j is given by (71), and iv) the government budget constraint (69) is satisfied.

Proposition. *Let*

$$t_L^k = \frac{1}{1+\nu}$$

$$\Delta_n^k = \sum_{k'j} w_n^{k'} \frac{L_n^{k'j}}{L_n^k} \frac{\partial \ln \lambda_n^{k'j}\left(\mathbf{L}_n\right)}{\partial \ln L_n^k}$$

and R^k such that

$$\phi^k U'(v^k) v^k L^k = \sum_n P_n C_n^k L_n^k.$$

Then, if the planner's problem is globally concave, the equilibrium with taxes and transfers implements the optimal allocation.

Proof. 1) The first order condition for household consumption choice (15) is identical to the first order condition for consumption in the planner's problem, (53). The modified budget constraint for the household (68) implies a relationship between consumption and prices identical to equation (64), which is derived from the first order conditions (55) and (56) in the planner's problem. At the same time, the optimal location decision for the household, (2) is identical to the free mobility constraint in the planner's choice (51) for a given set of consumption C_n^k.

2) The first order condition for factor demand for intermediate input producers, (17), (19) and (70) are identical to the first order conditions for the planner's problem (59), (60) and (61) once one uses equation (67) to substitute $\tilde{w}_n^k$ out of (59).

3) The condition that a producer in city n and industry j imports a variety $\mathbf{z}$ from city n' if and only if $\kappa_{nn'}^j p_n^j(\mathbf{z}) = \min_{n''} \kappa_{nn''}^j p_{n''}^j(\mathbf{z})$ is implied by the first order condition for the planner's problem (58), given that $Q_n^j(\mathbf{z}) = \sum_{n'} Q_{nn'}^j(\mathbf{z})$, $\tilde{p}_n^j(\mathbf{z})d\Phi(\mathbf{z}) = p_n^j(\mathbf{z})$.

4) The first order condition associated with the optimal use of different varieties by final goods producers (24) is implied by (58) given that $Q_n^j(\mathbf{z}) = \sum_{n'} Q_{nn'}^j(\mathbf{z})$, $\tilde{p}_n^j(\mathbf{z})d\Phi(\mathbf{z}) = p_n^j(\mathbf{z})$, and $P_n^j(\mathbf{z}) = \min_{n'} \kappa_{nn'}^j p_{n'}^j(\mathbf{z})$.

5) The market clearing conditions for employment (equation 5), structures (equation 6), final goods (equation 7) and intermediate goods (8) are identical to the resource constraints faced by the planner, respectively, (50) combined with (51), (52), (48) and (49).

6) In the planner's solution, equation 65 has to hold.

□

E A Counterfactual Economy After Eliminating Endogenous Amenities

We now verify whether the planner solution would be likely to change if one were to adjust local amenities to remove the components that Diamond (2016) argues are likely to be endogenous. For that purpose, we carry out two counterfactual exercises. For both exercises, we first extract the exogenous component of amenities as implied by the mapping of Diamond's (2016) estimates into amenity spillovers described in Fajgelbaum and Gaubert (2018). Specifically, we calculate a value of $A_n^{k,exo}$ such that $A_n^{k,exo} \prod \left(L_n^{k'}\right)^{\tau_a^{k'k}} C_n^k = 1$, with $\tau_a^{CNR,CNR} = 0.77$, $\tau_a^{nCNR,CNR} = -1.24$, $\tau_a^{CNR,nCNR} = 0.18$ and $\tau_a^{nCNR,nCNR} = -0.43$. In the first exercise, we calculate a counterfactual equilibrium where the labor supply equations are given by $L_n^k = \frac{\left(A_n^{k,exo} C_n^k\right)^{-\nu^k}}{\sum_{n'} \left(A_{n'}^{k,exo} C_{n'}^k\right)^{-\nu^k}}$. In the second exercise, we calculate the optimal allocation in that counterfactual environment.

Figures 21 below show the relationship between relative the exogenous part of amenities implied by that exercise and city size and composition, further discussed in the text.

Figure 22 shows how the distribution of CNR workers in the optimal allocation compares with the counterfactual equilibrium. As in our baseline economy, the planner has an incentive to increase labor market polarization by concentrating proportionately more CNR workers in larger cities. Figure 23 shows that, as in our baseline analysis, this increased polarization is matched by transfers from the large cities to the small ones.

Figure 21: Relative amenities and city size and composition (exogenous part)

Ratio of occupational-specific amenity parameters for each city obtained after extracting the endogenous part of amenities implied by the parametrization used by Fajgelbaum and Gaubert (2018). Each observation refers to a CBSA. Marker sizes are proportional to total city employment.

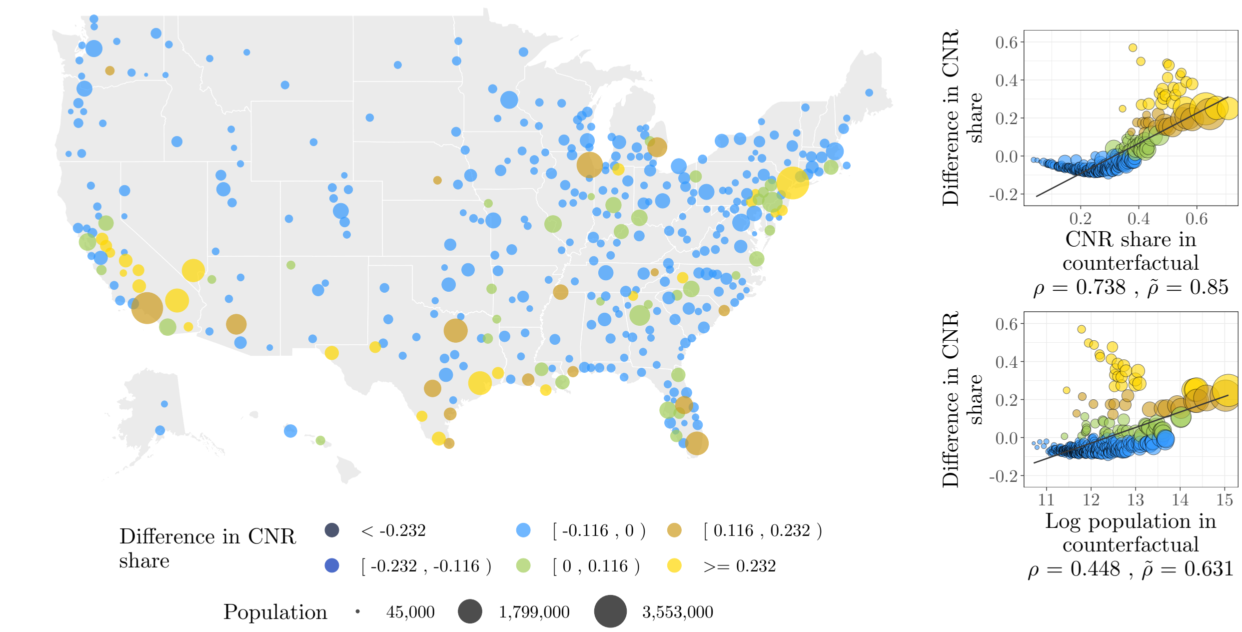

Figure 22: Optimal L_n^{CNR}/L_n with counterfactual amenities (change from counterfactual equilibrium)

Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

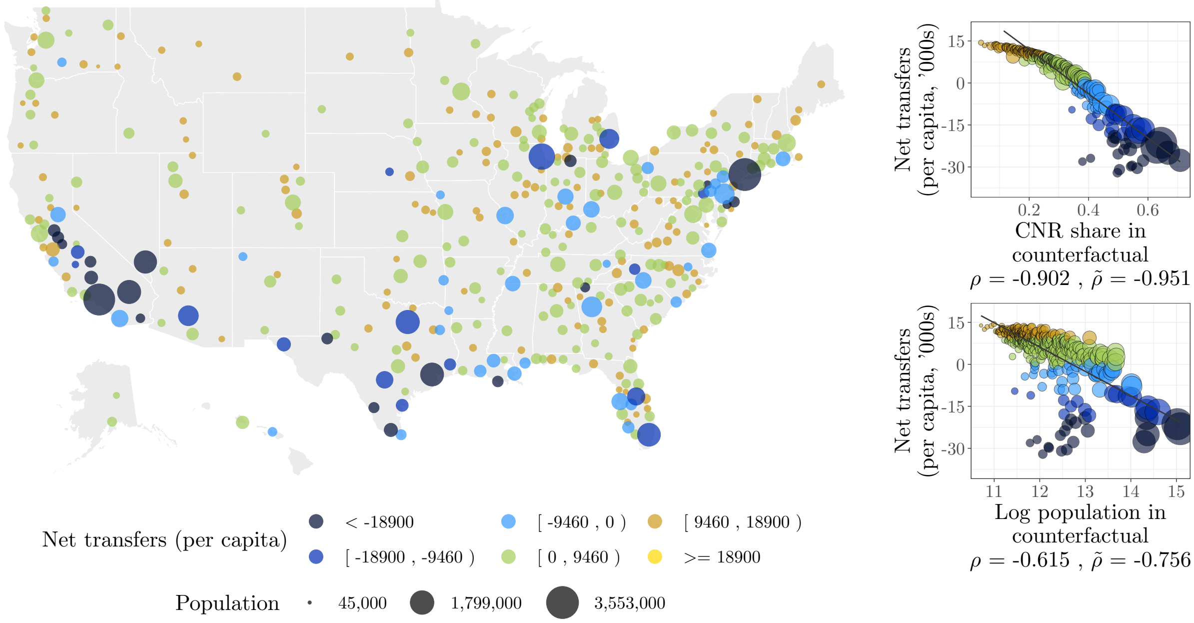

Figure 23: Optimal transfers with counterfactual amenities

Optimal transfers defined as the difference in the optimal allocation between the value consumed and value added in each city ($\sum_k P_n C_n^k - \sum_k w_n^k L_n^k - r_n H_n$). Each marker in the map refers to a CBSA. Marker sizes are proportional to total equilibrium employment in each city. ρ and $\tilde{\rho}$ are unweighted and population weighted correlations respectively.

F Quantifying the Model for 1980 and Counterfactual Exercises

F.1 Quantifying the Model for 1980

In order to quantify the model for 1980, we follow similar steps as described in Section B, with modifications to accommodate data constraints.

Regional Price Parities data are not available for 1980. In the baseline model quantification, we used those in order to calculate the productivity of the non-tradable sectors. To obtain the productivity of the real estate sector in 1980, we match instead changes in CoreLogic HPI data, available by county. As for the productivity of the non-tradable sector, we assume that its spatial distribution does not change. In addition, the model inversion exercise carried out for our 2011-15 benchmark does not pin down the national average level of productivity for each industry, only its occupational and spatial variation. In order to obtain the time variation of those levels, we choose average 1980 productivity levels to match national level sectoral price series made available by the BEA.

To obtain wages and the occupational composition of cities and industries, we use the 5% sample of the 1980 Census data which is comparable to the ACS. The 1980 Census has data for 213 MSA's that account for approximately 85% of U.S. employment in that year. For the remaining MSA's, we impute wages and employment by occupation and by sector by taking the predicted values of a regression of those variables on 1980 CBP employment by sector and housing prices.

F.2 Details of Counterfactual Exercises

In the counterfactual exercises described in Section 6, we separate average changes in productivity or amenities from their geographical and occupational dispersion.

The first step is to study the consequences of changing factor shares. We focus on the consequences of those changes to factor demand, while keeping unit costs in individual cities and industries fixed. This exercise implies a set of alternative productivity parameters for 1980, which we then take as our base for comparison with the current period.[39] Productivity changes then refer to changes in $T_n^{kj} = \left(H_n^j\right)^{\beta_n^j} \left(\lambda_n^{kj}\right)^{\gamma_n^j(1-\beta_n^j)}$.[40] The average change in productivity between 1980 and 2011-15 for a given industry is a Tornqvist type index: a geometric weighted average of the changes in productivity across cities, with the weights given by the value added by each city/industry as a fraction of total industry value added. Those shares are first calculated separately for the 1980 and 2011-15 periods, and the weights correspond to the arithmetic average of those shares.[41]

The model does not allow us to pin down an aggregate trend in amenities since changing amenities in all cities by a common scaling parameter leaves the equilibrium unchanged. We thus assume that there was no such trend so as to focus on the welfare implications of endogenous changes in equilibrium variables. For the baseline economy, this implies keeping a Tornqvist type index of amenities constant relative to the 2011-15 period: specifically, we keep a weighted geometric average of changes in amenities equal to 1, with the weights given by employment shares by city (again the shares are taken for the baseline and 2011-15 periods separately and the weights are given by an arithmetic average).

[39] One advantage of this procedure is that, given that changes in factor shares can be city-specific, implied productivity changes may otherwise depend on scaling parameters adopted for the different inputs.

[40] Specifically, when calculating the average change in productivity for a given sector j and occupation k we set $\ln\left(T_n^{kj,\text{counterfactual}}\right) = \gamma_n^j \sum_{n'} \omega_{n'}^{kj} \frac{1}{\gamma_{n'}^j} \ln\left(T_{n'}^{kj}\right)$, where $\omega_n^{kj} = \frac{w_n^k L_n^{kj}}{\sum_{n'} w_{n'}^k L_{n'}^{kj}}$, and analogously for other averages.

[41] We carry out a similar calculation in order to obtain productivity trends by city/industry/occupation

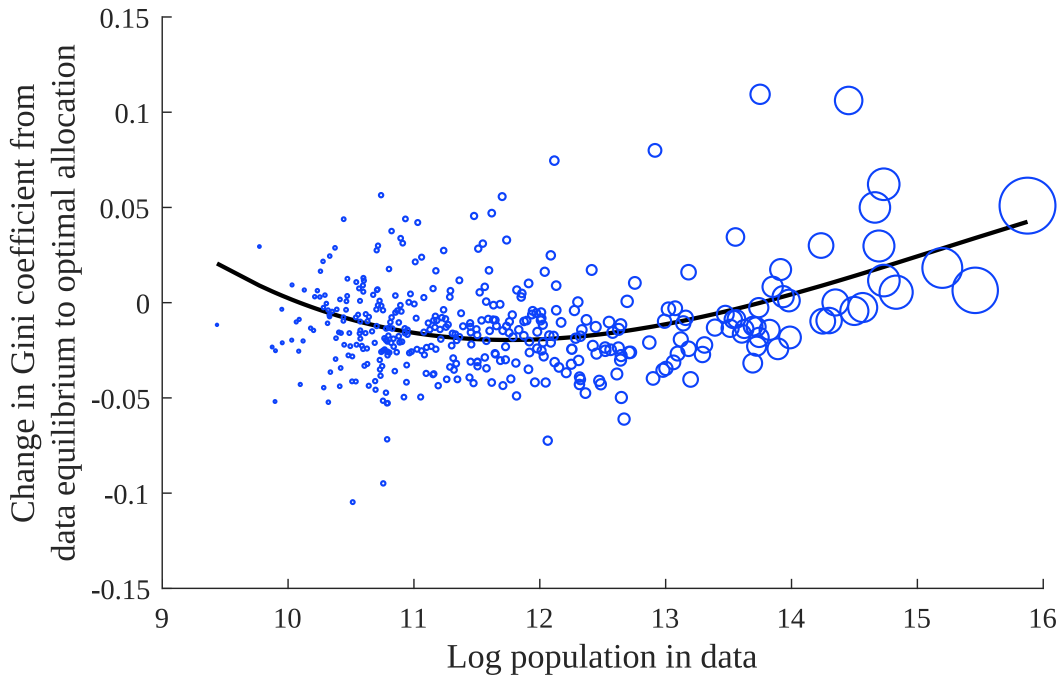

Figure 24: Changes in the Gini coefficient between the data and optimal allocation.

Each observation refers to a CBSA. Marker sizes are proportional to total employment. The solid-black line is a cubic fit of the data. The Gini is constructed using the Lorenz curves depicting within city wage bill and industry rank.