

Spatial Sorting

Jan Eeckhout

University College London, Barcelona Graduate School of Economics, and Universitat Pompeu Fabra

Roberto Pinheiro

University of Colorado

Kurt Schmidheiny

University of Basel

We investigate the role of skill complementarities in production and mobility across cities. The nature of the complementarities determines the equilibrium skill distribution across cities. With extreme-skill complementarity, the skill distribution has thicker tails in large cities; with top-skill complementarity, there is first-order stochastic dominance. Using wage and housing price data, we find robust evidence of thick tails in large cities: large cities disproportionately attract both high- and low-skilled workers, while average skills are constant across city size. This pattern of spatial sorting is consistent with extreme-skill complementarity, where the productivity of high-skilled workers and of the providers of low-skilled services are mutually enhanced.

If I can make it there I'll make it anywhere. (Frank Sinatra, "New York, New York")

Rock Bottom, yeah I see you, all my Detroit people. (Eminem, "Welcome 2 Detroit")

We are grateful to numerous colleagues for detailed discussion and insightful comments. We also benefited from the feedback of many seminar audiences. Eeckhout gratefully acknowledges support by the European Research Council, grant 208068, and Schmidheiny

[*Journal of Political Economy*, 2014, vol. 122, no. 3]

© 2014 by The University of Chicago. All rights reserved. 0022-3808/2014/12203-0005\$10.00

I. Introduction

Complementarities are important for the productivity and composition of workers in firms, for student achievement in classrooms (peer effects), and for the accomplishments of teams. The presence of more productive coworkers affects the performance of some, if not all, other coworkers, and this in turn determines who chooses to work where and with whom. In this paper we investigate the role of complementarities at an aggregate level—the level of a city. Complementarities are akin to knowledge spillovers (Marshall 1890), but rather than flowing between innovating firms, complementarities affect the productivity of differentially skilled workers within the local labor market. We propose a model that elucidates both the nature of cities and the role of complementarities in production. Our main theoretical finding is that the specifics of the complementarities determine the distribution of skills within a city and how it varies with city size. Our approach sheds new light on the sources of the urban wage premium, a major puzzle in the literature. It is well known that wages in large cities are higher, but it is unclear why. Little is known about the skill composition across cities. Are wages higher because workers in large cities are more skilled? Most people can provide casual evidence that the skill level in the top percentiles of New York and large cities in general is higher than anywhere else. Making it there—in New York—rather than in Akron, Ohio, is the ultimate aim of many professionals in many trades: artists, musicians, advertising and media professionals, consultants, lawyers, financiers, and so forth. Yet, cities are not just populated by superstars and high-earning professionals, even if these are highly visible.

In this paper, we address the sorting decision of workers over the entire range of skills, including medium and low skills. Our main empirical and hitherto undocumented finding is that the distribution of skills in the United States has thick tails in large cities: large cities disproportionately attract both high- and low-skilled workers, while average skills are constant across size. From the theory, this allows us to conclude that there are complementarities between high- and low-skilled workers, which mutually boosts their productivity.

We consider two competing hypotheses concerning the complementarities between skills. A first hypothesis is that the superstars boost their productivity most in the presence of other high-skilled workers. For example, under this assumption, the best lawyers are more productive when surrounded by top legal assistants. Or the cancer surgeons at Sloan Kettering in New York work best with top residents and top nurses, whereas

by the Swiss National Science Foundation, grant Sinergia/130648. Data are provided as supplementary material online.

the general practitioners with fewer years of training and fewer fellowships collaborate with less trained nurses and assistants. We refer to this hypothesis as *top-skill complementarity*. A second hypothesis posits that high-skilled workers boost their productivity most with low-skilled services. What dominates in the aggregate is that the high-skilled worker has a disproportionately high productivity increase from the presence of low-skilled services. Given the value of her time, at her job she hires more low-skilled administrative help and other services (sales, legal, catering, etc.) and she demands low-skilled services through child care, schooling, and help in the household. We label this hypothesis *extreme-skill complementarity*.

The premise of our analysis is that the presence of those complementarities determines the location decision of differentially skilled workers, that is, spatial sorting. We propose a theory that identifies a one-to-one relationship between those features of the technology, on the one hand, and the equilibrium outcome of the skill distributions across cities, on the other. Our model is a tractable version of the multiworker matching model à la Kelso and Crawford (1982) applied to a concrete labor market setting. Complementarities determine competitive wages and therefore the location decision of workers. Our objective is to uncover the nature of the complementarities from the observed sorting pattern of workers, that is, the skill distribution across cities. This is very much in the spirit of Krusell et al. (2000) in the macro literature, who derive properties of complementarities in technology from the observed wage distributions.

Our labor market model is tailored to investigate the nature of cities, and there are three contributions of the paper: First, we identify a mechanism of skill complementarities and the resulting skill distribution that we can explicitly solve, despite the fact that models with varying elasticity of substitution are notoriously hard to solve analytically, as Krusell et al.'s (2000) dynamic model illustrates. Second, qualitatively we discover an extremely robust empirical pattern of thick tails in the distribution of skills: average skills are independent of city size, while the standard deviation increases with city size. In conjunction with the results from the theory, this allows us to conclude that the observed pattern of skills is due to the complementarities between extreme skills. We believe that the theoretical link between the complementarities and the distribution in local labor markets is both theoretically and empirically novel. Third, our analysis makes further headway in our understanding of one of the major outstanding puzzles in urban economics, namely, the mechanism behind the urban wage premium. Our findings establish that wages are not higher because skills are uniformly higher. In fact, average skills are constant across cities. Wages are therefore higher only to compensate for higher housing prices. But our findings also show that the skill composition is important for the distribution of wages and productivity. Complementarities between extreme skills act as a multiplier of existing differences in total factor pro-

ductivity (TFP) across cities. Our theory of differential complementarities provides an explanation for higher wage and skill inequality in large cities.

The normative implications in the baseline model are particularly relevant when evaluating inequality. Our results show that from a social welfare viewpoint, wage inequality and urbanization are intimately related. This can have far-reaching policy implications. Consider, for example, the current income taxation system that progressively taxes individuals and households on their nominal income. Effectively, given the urban wage premium, this means that the current fiscal system differentially taxes cities of different sizes. The implication is that the population distribution across different-size cities is distorted, and as a result, aggregated output produced is suboptimal.

Prices play a key role in our equilibrium model of city choice. Heterogeneously skilled citizens earn a living based on a competitive wage and choose housing in a competitive housing market. Under perfect mobility, their location choice will make them indifferent between consumption-housing bundles and therefore between different wage-housing price pairs across cities. Wages are generated by firms that compete for labor and that have access to a city-specific technology summarized by that city's TFP. This naturally gives rise to a *price-theoretic measure of skills*. Larger cities pay higher wages and are more expensive to live in. Under worker mobility, revealed preference location choices imply that wages adjusted for housing prices are a measure of skills.

Using this price-based measure of skills, we can establish two robust empirical facts: average skills are constant across cities, and the standard deviation increases with city size. Big cities are characterized by big real inequality. The city size–wage premium is thus not driven by a high average skill level. Instead, larger cities have thicker tails in the skill distribution and disproportionately attract both higher- and lower-skilled agents. In New York City, for example, not only is there a huge contingent of high-skilled workers in Manhattan, but there are also disproportionately many low-skilled workers living in the South Bronx and Newark, New Jersey. Similarly, while Detroit has disproportionately many low-skilled individuals and a reputation for inner-city poverty, it also disproportionately attracts high-skilled individuals, many of whom live in the wealthy neighborhood of Bloomfield Hills. In that respect, large cities such as New York and Detroit are more similar to each other than they are to small cities because of the systematic pattern of thick tails in the skill distribution of large cities.

We document that this systematic pattern of spatial sorting is extremely robust to different measures: we use educational attainment and occupation as direct measures of skills and control for industry selection, we investigate the role of migration, we consider different definitions of large versus small cities, we use three different data sources for local hous-

ing values, and we include local price differences in consumption goods. To our knowledge, this pattern of spatial sorting—that mobility across cities driven by differential skill complementarity determines the skill distribution—has not been documented in the literature.

The thick tails in the skill distribution are nonetheless consistent with the well-documented city size–wage premium. The gap between average wages in the smallest cities in our sample (with a population of around 160,000, less than 1 percent of the population of New York) and the largest cities is 25 percent. In the left panel of figure 1, we plot a kernel of the wage distribution of those living in all cities larger than 2.5 million inhabitants and that of those in cities smaller than 1 million inhabitants. Not only are average wages higher, but there is a clear first-order stochastic dominance relation. At all wage levels, more people earn less in small cities than in large cities. This clearly indicates that there is a city size–wage premium across the board.

However, larger cities tend to be more expensive to live in, so in order to be able to compare skill distributions, we need to adjust for housing prices. Identical agents will make a location choice on the basis of the utility obtained, which depends on both wages and the cost of housing. The location decision will therefore require indifference for identical agents. We use homothetic preferences to adjust for housing consumption and construct a housing price index based on a hedonic regression to calculate the difference in housing values across cities. The right panel of figure 1 displays the kernel of the induced skill distribution. Our main finding is that the skill distribution in larger cities has fatter tails both at the top and at the bottom of the distribution. Large cities disproportionately attract more skilled and more unskilled workers. This finding sheds light on the nature of the underlying technology: the fat tails result is consistent with extreme-skill complementarity.

A key feature of our approach is the price-theoretic measure of skills that allows us to characterize a smooth distribution of skills. This is in contrast to the common approach of using observable skills such as attained education levels or years of education. To investigate the role of observables in the spatial sorting pattern, we decompose the difference in the skill distribution between large and small cities. We find an asymmetry: in the lower tail, virtually none of the city size difference is explained by observables, while in the upper tail about half is explained, mainly by education. The high-skilled are more educated in large cities than in small cities, while the low-skilled are equally educated across city sizes.

There are of course other possible alternative explanations. And while we cannot exhaustively analyze all alternatives empirically, we can rule out a few prominent candidate alternatives and establish the robustness of our findings with respect to industry composition, migration, and age

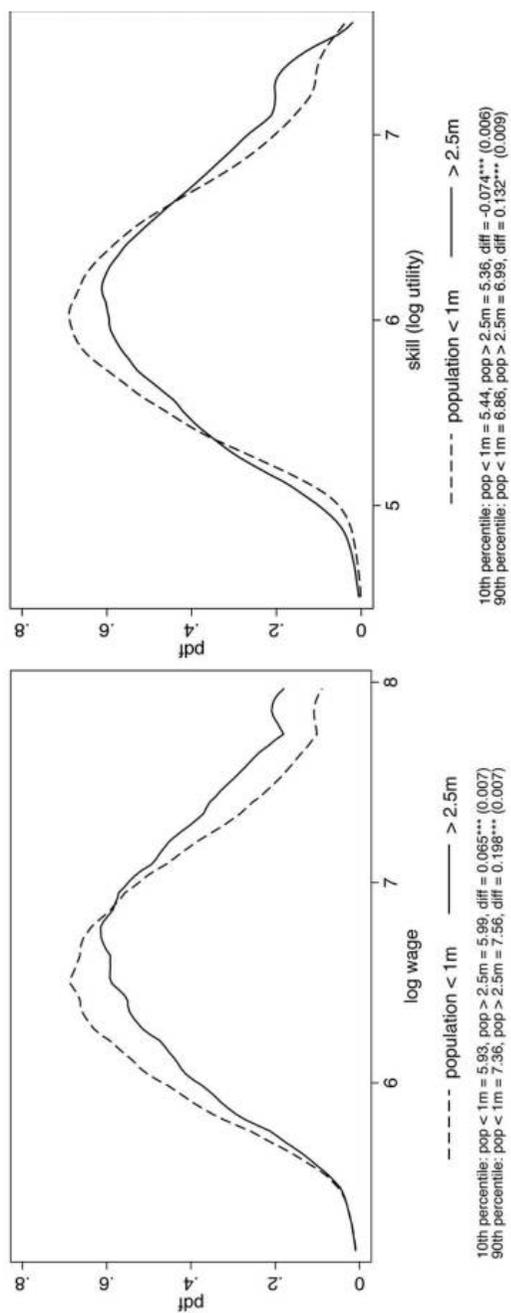


FIG. 1.—Wage and skill distribution for small and large cities, kernel density estimates (Epanechnikov kernel, bandwidth = 0.1); wage data from the 2009 CPS on 25,584 workers in 202 small CBSAs (population between 100,000 and 1 million) and 34,999 workers in 21 large CBSAs (larger than 2.5 million): *left*, wages; *right*, skills. Standard errors are in parentheses; * $p < .10$, ** $p < .05$, *** $p < .01$.

or life cycle patterns. We discuss competing theoretical explanations such as the role of home versus market production in a world in which agents have preferences for low-skilled services. We also investigate the role of nonhomothetic preferences and within-city sorting.

II. Related Literature

The model we propose builds on the urban location model in Eeckhout (2004) and Davis and Ortalo-Magné (2011) (see also Guerrieri, Hartley, and Hurst [2011], who augment the model with local externalities) in which identical citizens who have preferences over consumption and housing choose a city in order to maximize utility. This model has been used to explain population dynamics (see also Gabaix 1999) and expenditure shares. Because of differences in productivity across cities, wages differ and housing prices adjust as a function of the population size of the city. Productivity differences are due to TFP and agglomeration effects. Given perfect mobility and identical agents, utility equalizes across cities. Our main innovation over the existing model is the introduction of heterogeneity in the inputs of production (skills), which gives rise to a distribution of skills within the city. This is necessary to meaningfully address sorting of heterogeneous agents within and across cities. Technology allows for varying degrees of complementarities between different skill types. Equilibrium is determined by the sorting decision of agents.

In recent work, Behrens, Duranton, and Robert-Nicoud (2014; in this issue) analyze the distribution of heterogeneous agents across cities. Their model predicts perfect sorting by talent: New York attracts all the PhDs, Los Angeles and Chicago all the masters, and so forth, and all the high school dropouts locate in small cities like Janesville, Wisconsin. Within-city heterogeneity in productivity is due to an *ex post* shock on which workers cannot relocate any more. As a result, they postulate first-order stochastic dominance of the (degenerate) talent distributions rather than the thick tails that our model predicts and that we find in the data. Their result of perfect sorting in talent is a direct consequence of assuming that each worker consumes one unit of land independent of his wage. Highly talented workers with high wages are therefore, relatively, much less affected by high housing costs in large cities than less talented workers. In contrast, we do allow housing consumption to increase with wages and hence with talent in line with overwhelming empirical evidence. In the equilibrium allocation of our model, all skill types locate in all cities simultaneously, driven by complementarities in production. This not only gives rise to a wage and skill distribution with full support as observed across cities of all sizes but, most importantly, also allows us to infer the pattern of complementarities in the technology that drive the location decision of workers.

Finally, Van Nieuwerburgh and Weill (2010) set up a spatial equilibrium model to explain the increase in housing price dispersion resulting from an increase in productivity dispersion of heterogeneous workers. As in our model, worker mobility in response to productivity shocks and endogenous housing prices are the main ingredients of their explanation. The key difference is the production function for consumption goods, which is linear in labor and does not feature skill complementarities.

There is a long tradition in the urban economics literature investigating differences across city sizes, in particular with respect to varying standards of living between cities. Albouy (2008) calculates real urban wages for 290 metropolitan statistical areas (MSAs) using the 2000 census (5 percent Integrated Public Use Microdata Series [IPUMS]). Nominal wages are deflated using rental prices from the census and local prices for consumption goods. The American Chamber of Commerce Research Association (ACCRA) Cost of Living Index is the basis of the latter but is not directly used because of its limited quality. Albouy regresses the ACCRA index on local rental prices and uses the predicted values as an index for local cost-of-living differences. Differences in real wages across MSAs are interpreted as quality-of-life differences. He finds that when local differences in federal taxes, nonlabor income, and observable amenities such as seasons, sunshine, and coastal location are controlled for, quality of life does not depend on size.

This body of work is consistent with our finding that the average of the skill distribution is remarkably constant across different-sized cities. Of course, that does not allow us to conclude that there is no sorting of high-skilled workers into large cities and of low-skilled workers into small cities. As we will show below, the data reveal quite the contrary. The mean is constant across cities of different size, but the variance increases substantially. The latter indicates an important role for sorting of high- and low-skill types into large cities and of medium types into small cities.

Our findings are also related to the previous literature on variations in the measured skill distributions across city sizes. Bacolod, Blum, and Strange (2009) study the difference in skill distributions across city sizes using jointly census and National Longitudinal Study of Youth data and the Dictionary of Occupational Titles, defining skills as a combination of qualities instead of just education. They find a small variation in cognitive, people, and motor skills across city sizes that they attribute to skills being defined nationally. However, they are not able to address local differences in occupational requirements of skills. Once they look at differences in the Armed Forces Qualification Test and the Rotter index—measures of intelligence and social skills, respectively—they find that, even though the average scores are quite similar across city sizes, the lowest scores in large cities (10th percentile) are much lower than the ones in small cities. Similarly, the highest scores (90th percentile) are much higher in large cit-

ies than in small ones, which is consistent with our robustness exercise reported in Section VII.A on direct measures of education. Also, Gautier and Teulings (2009, table 1) report a higher measured standard deviation in educational attainment across larger cities, which is consistent with our findings. However, they find first-order stochastic dominance rather than thick tails as we do because the mean is also higher. The reason is that their measure of skill is really a measure of wages. It is independent of housing prices, and it is constructed as predicted wages net of unobserved heterogeneity using a Mincerian wage regression. Consistent with the urban wage premium, average wages increase with city size. Instead, our measure of skills adjusts wages for the equilibrium mobility decision by means of housing prices, and we find that average skills are independent of city size. Together with the fact that the standard deviation increases with city size, this gives us thick tails and not first-order stochastic dominance. Note also that the first-order stochastic dominance in Gautier and Teulings' study is not consistent with the direct measures of skills as reported in Section VII.A or as documented by Bacolod et al. (2009). The distributions of those direct measures have thick tails in large cities—equal means, higher standard deviation—just as our wage-based measure that is adjusted for housing prices. In sum, while the wage distribution has been shown to satisfy first-order stochastic dominance in city size (increasing mean, increasing standard deviation), we establish that the distribution of our wage-based measure of skills has thick tails in large cities (equal mean, increasing standard deviation).

There is also recent literature on increasing wage inequality over time (Baum-Snow and Pavan 2012; Autor and Dorn 2013; Moretti 2013). Autor and Dorn document faster growth at both tails of the wage distribution between 1980 and 2005 and attribute this to the falling cost of automating middle-skill routines, the polarization hypothesis.

Finally, there is little direct evidence on the role of complementarities between heterogeneously skilled agents. One notable exception is the work by Hamilton, Nickerson, and Owan (2003), who analyze the effect of team composition on productivity in a textiles production plant. They find that heterogeneous teams are more productive with average productivity held constant. While their setup is very specific and other theories can certainly rationalize this outcome, their finding is consistent with a technology that has extreme-skill complementarity.

III. The Model

Population.—Consider an economy with heterogeneously skilled workers. Workers are indexed by a discrete skill type $i: i \in \mathcal{I} = \{1, \dots, I\}$. Associated with this skill order is a level of productivity y_i . Denote the countrywide measure of skills of type i by M_i . Let there be J locations (cities)

$j \in \mathcal{J} = \{1, \dots, J\}$. The amount of land in a city is fixed and denoted by H_j . Land is a scarce resource.

Preferences.—Citizens of skill type i who live in city j have preferences over consumption c_{ij} and the amount of land (or housing) h_{ij} . The consumption good is a tradable numeraire good with price equal to one. The price per unit of land is denoted by p_j . We think of the expenditure on housing as the flow value that compensates for the depreciation, interest on capital, and so forth. In a competitive rental market, the flow payment will equal the rental price.¹ A worker has consumer preferences over the quantities of goods and housing c and h that are represented by $u(c, h) = c^{1-\alpha} h^\alpha$, where $\alpha \in [0, 1]$. Workers are perfectly mobile, so they can relocate to another city instantaneously and at no cost. Because workers with the same skill are identical, in equilibrium each of them should obtain the same utility level wherever they choose to locate. Therefore, for any two cities j, j' , it must be the case that the respective consumption bundles satisfy $u(c_{ij}, h_{ij}) = u(c_{ij'}, h_{ij'})$ for all skill types for all $i \in \{1, \dots, I\}$.

Technology.—Cities differ in their TFP, which is denoted by A_j . For now, we assume that TFP is exogenous. We think of it as representing a city's productive amenities, infrastructure, historical industries, persistence of investments, and so forth.²

In each city, there is a technology operated by a representative firm that has access to a city-specific TFP A_j . Output is produced by choosing the right mix of differently skilled workers i . For each skill i , a firm in city j chooses a level of employment m_{ij} and produces output $A_j F(m_{1j}, \dots, m_{Ij})$. Firms pay wages w_{ij} for workers of type i . It is important to note that wages depend on the city j because citizens freely locate between cities not on the basis of the highest wage but, given housing price differences, of the highest utility. Like land, firms are owned by absentee capitalists (or, equivalently, all citizens own an equal share in the mutual fund that owns all the land and all the firms).

Market clearing.—In the countrywide market for skilled labor, markets for skills clear market by market, and for housing, there is market clearing within each city:

$$\begin{aligned} \sum_{j=1}^J C_j m_{ij} &= M_i \quad \forall i, \\ \sum_{i=1}^I h_{ij} m_{ij} &= H_j \quad \forall j, \end{aligned} \tag{1}$$

where C_j denotes the number of cities with TFP A_j .

¹ We will abstract from the housing production technology; e.g., we can assume that the entire housing stock is held by a zero measure of absentee landlords.

² In an earlier version of the paper, we endogenize A_j and let it be the result of agglomeration externalities. This is also documented in the online supplementary appendix.

IV. The Equilibrium Allocation

The citizen's problem.—Within a given city j and given a wage schedule w_{ij} , a citizen chooses consumption bundles $\{c_{ij}, h_{ij}\}$ to maximize utility subject to the budget constraint (where the tradable consumption good is the numeraire, i.e., with price unity):

$$\begin{aligned} \max_{\{c_{ij}, h_{ij}\}} u(c_{ij}, h_{ij}) &= c_{ij}^{1-\alpha} h_{ij}^{\alpha} \\ \text{subject to } c_{ij} + p_j h_{ij} &\leq w_{ij} \end{aligned} \quad (2)$$

for all i, j . Solving for the competitive equilibrium allocation for this problem, we obtain $c_{ij}^* = (1 - \alpha)w_{ij}$ and $h_{ij}^* = \alpha(w_{ij}/p_j)$. Substituting the equilibrium values in the utility function, we can write the indirect utility for a type i as

$$\begin{aligned} U_i &= \alpha^{\alpha} (1 - \alpha)^{1-\alpha} \frac{w_{ij}}{p_j^{\alpha}} \\ \Rightarrow w_{ij} &= U_i p_j^{\alpha} \frac{1}{\alpha^{\alpha} (1 - \alpha)^{1-\alpha}}, \end{aligned} \quad (3)$$

where U_i is constant across cities from labor mobility. This allows us to link the wage distribution across different cities j, j' . Wages across cities relate as

$$\frac{w_{ij}}{w_{ij'}} = \left(\frac{p_j}{p_{j'}} \right)^{\alpha}. \quad (4)$$

The firm's problem.—All firms are price takers and do not affect wages. Wages are determined simultaneously in each submarket i, j . Given the city production technology, a firm's problem is given by

$$\max_{m_{ij} \forall i} A_j F(m_{1j}, \dots, m_{ij}) - \sum_{i=1}^I w_{ij} m_{ij} \quad (5)$$

subject to the constraint that $m_{ij} \geq 0$. The first-order condition is $A_j F_{m_{ij}}(m_{ij}) = w_{ij}$ for all i .³

Because there is no general solution for the equilibrium allocation in the presence of an unrestricted technology, we focus on variations of the

³ In what follows, the nonnegativity constraint on m_{ij} is dropped. This is justified whenever the technology satisfies the Inada condition that marginal product at zero tends to infinity whenever A_j is positive. This will be the case since we focus on variations of the constant elasticity of substitution (CES) technology.

CES technology, where the elasticity is allowed to vary across skill types. As a benchmark, therefore, we consider the CES technology

$$A_j F(m_{1j}, \dots, m_{lj}) = A_j \left(\sum_{i=1}^l m_{ij}^\gamma y_i \right) \tag{6}$$

with $\gamma < 1$. In this case, the first-order conditions are $A_j \gamma m_{ij}^{\gamma-1} y_i = w_{ij}$ for all i .

Below we will solve the allocation under CES as a special case of the more general technology. Even without fully solving the system of equations for the equilibrium wages, observation of the first-order condition reveals that productivity between different skills i in a given city is governed by two components: (1) the productivity y_i of the skilled labor and how fast it increases in i and (2) the measure of skills m_{ij} employed (wages decrease in the measure employed from the concavity of the technology). Without loss of generality, we assume that wages are monotonic in the order i .⁴ This is consistent with our price-theoretic measure of skill.

We now proceed by introducing varying degrees of complementarities between different skills, starting from the CES technology. This implies that the technology now has an elasticity of substitution that is no longer constant. For tractability, let there be two cities, $j \in \{1, 2\}$, and three skill levels, $i \in \{1, 2, 3\}$. Consider any subset of the skills, say i, k , between which there is a degree of complementarity λ , and none with the remaining skill level l . Then the technology can be written as $(m_{ij}^\gamma y_i + m_{kj}^\gamma y_k)^\lambda + m_{lj}^\gamma y_l$. Depending on the subset of skills, we distinguish between the following configurations.

DEFINITION 1. Consider the following technologies:

- I. Extreme-skill complementarity: High-skill workers are complementary with low-skill workers,

$$A_j F(m_1, m_2, m_3) = A_j [(m_{1j}^\gamma y_1 + m_{3j}^\gamma y_3)^\lambda + m_{2j}^\gamma y_2], \tag{7}$$

when $\lambda > 1$ relative to CES. Instead, skills 1 and 3 are substitutes when $\lambda < 1$.

- II. Top-skill complementarity: High-skill workers are complementary with medium-skill workers,

$$A_j F(m_1, m_2, m_3) = A_j [(m_{2j}^\gamma y_2 + m_{3j}^\gamma y_3)^\lambda + m_{1j}^\gamma y_1], \tag{8}$$

⁴ For a given order i , wages may not be monotonic as they depend on the relative supply of skills as well as on y_i . If they are not, we can relabel skills such that the order i corresponds to the order of wages. Alternatively, we can allow for the possibility that higher-skilled workers can perform lower-skilled jobs. Workers will drop job type until wages are nondecreasing. Then the distribution of workers is endogenous, and given this endogenous distribution, all our results go through. For clarity of the exposition, we will assume that the distribution of skills ensures that wages are monotonic.

when $\lambda > 1$ relative to CES. Instead, skills 2 and 3 are substitutes when $\lambda < 1$.

Observe that we could also introduce bottom-skill complementarities. In terms of the distributional implications, this is equivalent to top-skill substitutabilities, that is, technology II with $\lambda < 1$. There are therefore five distinct configurations of the technology: two for technology I, with complements ($\lambda > 1$) or substitutes ($\lambda < 1$), two for technology II ($\lambda > 1$ and $\lambda < 1$), and CES ($\lambda = 1$).

It is worth pointing out that for our purpose, three skills is the minimal requirement to fully capture first-order stochastic dominance and thick tails. Distinguishing between the two cannot be achieved with two skills only. At the same time, with a larger number of skills, we do not obtain qualitatively different results. With 100 skill types, one can of course analyze the properties of each percentile, but that does not provide essential additional information about the existence of thick tails or stochastic dominance. We nonetheless investigate the generality of this setup. In the online supplementary appendix, we report the same properties that we derive below for general technologies with any N skills and for more general patterns of gross complementarities. For expositional purposes, in what follows we focus on the case of two city types A_1, A_2 with any C_1, C_2 , but we can also analyze the general case with J cities.

We first derive the equilibrium conditions for case I, extreme-skill complementarity. The first-order conditions are for each j and all skill types i , respectively:

$$\lambda A_j (m_{1j}^\gamma y_1 + m_{3j}^\gamma y_3)^{\lambda-1} \gamma m_{1j}^{\gamma-1} y_1 - w_{1j} = 0, \quad (9)$$

$$\gamma A_j m_{2j}^{\gamma-1} y_2 - w_{2j} = 0, \quad (10)$$

$$\lambda A_j (m_{1j}^\gamma y_1 + m_{3j}^\gamma y_3)^{\lambda-1} \gamma m_{3j}^{\gamma-1} y_3 - w_{3j} = 0. \quad (11)$$

Using labor mobility, we can write the wage ratio in terms of the house price ratio for all i , $w_{2i}/w_{1i} = (p_2/p_1)^\alpha$, and equate the first-order condition in both cities for a given skill, for example, for $i = 1$:

$$A_1 (m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{11}^{\gamma-1} = \left(\frac{p_1}{p_2} \right)^\alpha A_2 (m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{12}^{\gamma-1}. \quad (12)$$

Using market clearing, $m_{12} = (M_1/C_2) - (C_1/C_2)m_{11}$ in the local labor market, we can solve for the first-order conditions for each skill to obtain the equilibrium quantities:

$$\begin{aligned}
 m_{11} &= \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)} \frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}, \\
 m_{21} &= \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\gamma-1)} \frac{M_2}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\gamma-1)}}, \\
 m_{31} &= \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)} \frac{M_3}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}.
 \end{aligned}
 \tag{13}$$

and likewise in city 2.

So far we have consumer optimization for consumption and housing, the location choice by the worker, and firm optimization given wages. The next step is to allow for market clearing in the housing market given land prices. The system is static and is solved simultaneously, which is reported in the Appendix A. In what follows, we assume $H_j = H$ for all cities j . Below, we will discuss the implications in which this simplifying assumption has bite.

The main theoretical results.—First we establish the relation between TFP and city size. Denote by S_j the size of city j where $S_j = \sum_{i=1}^I C_j m_{ij}$. When cities have the same amount of land, we can establish the following result.

THEOREM 1 (City size and TFP). Let $A_1 > A_2$, $\lambda\gamma < 1$, and $\gamma < 1$. Then the more productive city is larger, $S_1 > S_2$.

Proof. In Appendix A.

We establish this result for cities with an identical supply of land. Clearly, the supply of land is important in our model since in a city with an extremely small geographical area, labor demand would drive up housing prices all else equal. This may therefore make it more expensive to live in even if the productivity is lower. Because in our empirical application we consider large metropolitan areas (New York City, e.g., includes large parts of New Jersey and Connecticut), we believe that this assumption does not lead to much loss of generality.⁵

⁵ In fact, the equal supply of housing condition is only sufficient for the proof, but not necessary. However, our model does not address the important issue of within-city geographical heterogeneity, as analyzed, e.g., in Lucas and Rossi-Hansberg (2002). In our application, all heterogeneity is absorbed in the pricing index by means of the hedonic regression. In Sec. VI.A, we empirically analyze the implications of within-city sorting and find no qualitative impact on the results. This is consistent with recent work by Fu and Ross (2010), who find little evidence of sorting within metropolitan areas based on agglomeration.

We now establish the main theorem characterizing the skill distribution across firms. We already know that more productive cities are larger, but this does not necessarily mean that the distribution of skills in larger cities differs from that in smaller cities. In fact, it depends on the technology.

THEOREM 2 (Extreme-skill complementarity and thick tails). Given $A_1 > A_2$, $\lambda > 1$, and $\lambda\gamma < 1$, the skill distribution in the larger city has thicker tails.

Proof. In Appendix A.

Two corollaries immediately follow from the main theorem.

COROLLARY 1 (CES technology). If $\lambda = 1$ and $\gamma < 1$, then the skill distributions across cities are identical.

Under CES technology, cities have identical skill compositions. This is due to the homotheticity of CES technology: the marginal rate of technical substitution is proportional to total employment, and, as a result, firms in different cities and with different technologies will employ different skills in the same proportions. The proof of the result follows immediately from setting $\lambda = 1$ in the proof of theorem 2. Even though the city skill distribution under CES technology is the same across cities, the more productive city will be larger. This follows from theorem 1.

The second corollary establishes the mirror-image result under extreme-skill substitutability.

COROLLARY 2 (Extreme-skill substitutability and thin tails). Given $A_1 > A_2$, $\lambda < 1$, and $\lambda\gamma < 1$, the skill distribution in the larger city has thinner tails.

These two corollaries can help build intuition for the result in theorem 2. Consider first CES as a benchmark. Homotheticity implies that even though the level of employment differs across skills, firms will always choose to hire different skills in exactly the same proportions for a given wage ratio. Since housing prices affect all skills within a city in the same way, the wage ratio is unaffected.

Instead with extreme-skill complementarity, the marginal product of both the low- and the high-skilled workers is higher than for medium skills, thus breaking the homotheticity. Given the complementarity between TFP A_j and the skill aggregator, the marginal impact on productivity of the extreme skills will now be disproportionately higher in larger than in smaller cities. This induces the relative increase in demand for extreme skills. Observe that this cannot be offset by higher housing prices because these are determined by real wage equalization at all skill levels, including the medium-skilled. The higher real wages for low- and high-skilled workers in large cities will attract those skill types into the large cities, driving down nominal wages until real wages are equalized. This immigration of low- and high-skilled workers leads to the thick tails in the large cities.

Top-skill complementarity.—Now consider the technology

$$A_j F(m_1, m_2, m_3) = A_j [(m_2^\gamma y_2 + m_3^\gamma y_3)^\lambda + m_1^\gamma y_1].$$

Without going through the detailed analysis in the text, we obtain the equivalent to theorem 2 above (theorem 1 readily extends as well).

THEOREM 3 (Top-skill complementarity and first-order stochastic dominance). Given $A_1 > A_2$, $\lambda > 1$, and $\lambda\gamma < 1$, the skill distribution in the larger city first-order stochastically dominates.

Proof. In Appendix A.

The following corollary establishes the mirror-image result under extreme-skill substitutability.

COROLLARY 3 (Top-skill substitutability and first-order stochastic dominance). Given $A_1 > A_2$, $\lambda < 1$, and $\lambda\gamma < 1$, the skill distribution in the larger city is first-order stochastically dominated.

Under top-skill complementarity, the highest-skilled are complements with the next-highest skill types, thus generating disproportionately higher output in larger cities. This complementarity breaks the homotheticity property and leads to disproportionate demand in larger cities. Free mobility and real wage equalization across cities imply that the distribution in the larger city has disproportionately more of the top skill types. This induces first-order stochastic dominance.

In theorems 2 and 3 we identify a mechanism of skill complementarities in the production technology that generates a systematic pattern in the skill distribution. There is exactly one distribution pattern that corresponds to each of the five technology patterns (extreme-skill and top-skill, each with complements or substitutes, and CES). From the systematic pattern of thick tails in the distribution in large cities that we observe below, we can qualitatively deduce that this is due to the complementarities between extreme skills. As in the macro literature on differential complementarities (most notably Krusell et al. [2000]), we obtain information about the technology from the observed equilibrium distribution. We believe that the theoretical link between the complementarities and the distribution is novel. Moreover, from a theoretical viewpoint, we are able to explicitly analyze a tractable matching problem à la Kelso and Crawford (1982) with complementarities (peer effects) that are applied in a concrete labor market setting. Thus far, only general properties such as existence rather than explicit characterizations have been analyzed in these models.

We have chosen to model city difference by means of exogenously given TFP differences. In reality, there are reasons why the productivity of cities is endogenous. We report a model with endogenous agglomeration externalities in the online supplementary appendix. The main find-

ing is that agglomeration externalities can lead to asymmetric equilibria with cities of different sizes, even if they are ex ante identical. This occurs provided that the external effect is strong enough. We further show that once cities are heterogeneous, the thick tail results extend to this setting with endogenous externalities. And in addition to these production externalities, there could be consumption externalities from the presence of amenities. Unfortunately in our analysis, because we identify unobservable skills from wages, we cannot jointly determine skills and amenities from the same wages. While there is no doubt that amenities matter for citizens' location decisions, on the basis of evidence from Albouy (2008), there seems to be no systematic relation to city size, unlike the relation of skill composition to city size as we derive in our results.

In Appendix C, we also discuss an alternative explanation, namely, that thick tails are generated by the combination of top-skill complementarity and the preference for services in a model with home production and a market for services. The intuition behind this model generating thick tails is simple: Top-skill complementarity would attract highly skilled workers to large cities, while the demand for services generated by these highly paid workers would attract low-skill workers to large cities as well. However, as we show in Appendix C, under reasonable parameter values, this setup would generate thick tails only under very specific conditions. First, the income share of services must be almost as high as the share on housing expenses. Second, the entry cost in the service sector must be sufficiently high. Finally, the top-skill complementarity must not be too strong. These conditions are very specific, and it is not at all clear that they are supported by the facts.

Finally, we also report some further results on housing and consumption expenditure. It is immediate from our model that in large cities, citizens will spend more on housing, yet they will consume less of it.

PROPOSITION 1. Consider a general technology F . For a given skill i , expenditure on housing $p_j h_{ij}^*$ is higher in larger cities. The size of houses h_{ij}^* in larger cities is smaller.

Proof. From the consumer's problem, we have $p_j h_{ij} = \alpha w_{ij}$. Since $w_{i1} > w_{i2}$, we must have $p_1 h_{i1} > p_2 h_{i2}$ for all i . Similarly, from the same equality in the consumer's problem, we have $h_{ij} = \alpha w_{ij} / p_j$. Since

$$\frac{w_{i1}}{p_1} < \frac{w_{i2}}{p_2}, \quad (14)$$

it follows that $h_{i1} < h_{i2}$. QED

Then given homothetic preferences for consumption, we immediately get the following corollary.

COROLLARY 4. Expenditure on the consumption good is higher in larger cities.

Our model predicts that expenditure on both housing and consumption is higher in larger cities, though the equilibrium quantity of housing h_{ij}^* is lower. As cities become larger (or as the difference in TFP increases), at all skill levels total income increases and therefore total expenditure increases. Because housing prices increase as well, there will be substitution away from housing to the consumption good. As a result, inequality in consumption expenditure will increase.

V. The Empirical Evidence of Thick Tails

We use the one-to-one relation between skills and equilibrium utility to back out the skill distribution from observable variables. The worker's indirect utility in equilibrium is independent of the city, given perfect mobility, and assuming Cobb-Douglas preferences, it satisfies

$$U_i = \alpha^\alpha (1 - \alpha)^{1-\alpha} \frac{w_{ij}}{p_j^\alpha}, \quad (15)$$

where we need to observe the distribution of wages w_{ij} by city j , the housing price level p_j by city, and the budget share of housing α .

A. Data

The analysis is performed at the city level. We define a city as a core-based statistical area (CBSA), the most comprehensive functional definition of metropolitan areas published by the Office of Management and Budget in 2000. See appendix table B1 for the list of the largest and smallest cities and their 2009 population.

We use wage data from the Current Population Survey (CPS) for the year 2009. We observe weekly pretax earnings for 76,821 full-time workers in 254 US metropolitan areas. CPS wages are top-coded at around \$150,000, which we will take into account in the statistical analysis. In Appendix B, we provide detailed information on data source, sample restrictions, and variables.

Local housing price levels are estimated using the American Community Survey (ACS) for 2009. We observe monthly rents for 273,761 housing units in 533 CBSAs. The ACS reports the number of rooms, the age of the structure, and the number of units in the structure. With these data we estimate city-specific housing price indices using hedonic regressions. See Appendix B for details and a theoretical motivation of this approach.⁶

⁶ In an earlier version of this paper we show that our findings are robust to using other housing price data such as from the 2000 US Census, the National Association of Realtors, or the Council for Community and Economic Research (C2ER).

B. Wage Distribution

Figure 1 (left panel) in the introduction shows the distribution of weekly wages for full-time earners both in cities with a population of more than 2.5 million and in cities with a population between 100,000 and 1 million. We clearly see that wages in larger cities are higher and that the top tail of the distribution is substantially fatter in large cities.⁷ A simple t -test shows that wages in large cities are 13.3 percent higher than in small ones ($t = 27.8$, $p < .01$). Controlling for right-censoring from top-coding and weights in a censored (Tobit) regression leads to almost exactly the same comparison: $\Delta \log \text{wage} = 13.2$ percent (robust $t = 24.7$, $p < .01$). A look at the tails of the two distributions shows that the large cities have a thicker tail at the top and the small cities at the bottom. The 90th percentile for large cities is 7.56 compared to 7.36 for small cities ($\Delta = 0.198$, standard error [SE] = 0.007, $p < .01$). The 10th percentile for large cities is 5.99 compared to 5.93 for small cities ($\Delta = 0.065$, SE = 0.007, $p < .01$).⁸

The above partitioning of wages into a group of small cities and a group of large cities is arbitrary. We therefore perform quantile regressions of wages on city population size. The left panel in figure 2 shows the estimated regression lines for the 10th, the 25th, the 50th, the 75th, and the 90th percentiles. The right panel in figure 2 shows the slope coefficients for all quantiles. The slopes are all significantly above zero, which implies that the upper tail of the wage distribution increases with city size while the lower tail decreases. For the median (50th percentile), for example, the slope is 0.042 (SE = 0.002, $p < .01$): a doubling of city size leads to a 4.2 percent increase in wages.

C. Skill Distribution

Davis and Ortalo-Magné (2011) document that expenditure shares on housing are remarkably constant across US metropolitan areas with a median expenditure share of 0.24. We use this as our estimate of α . Together with our estimate for local housing prices p_j , we can back out the indirect utility u_{ij} for the observed wages using equation (15).

The variation in housing prices is substantial. While wages increase by 4.2 percent as city size doubles, housing prices increase by 16.9 percent for the same change in city size, that is, a fourfold increase. With the 0.24 expenditure share, this implies that the average cost of living is of a factor $1.169^{0.24} = 1.038$. In other words, the 4.2 percent wage gain from living

⁷ Note that the “bumps” in the top tail for both large and small cities are an artifact of the top-coded nominal wage data.

⁸ Percentiles and their difference are estimated in a quantile regression of wages on a dummy variable for large cities. We use CPS earnings weights and bootstrapped robust standard errors.

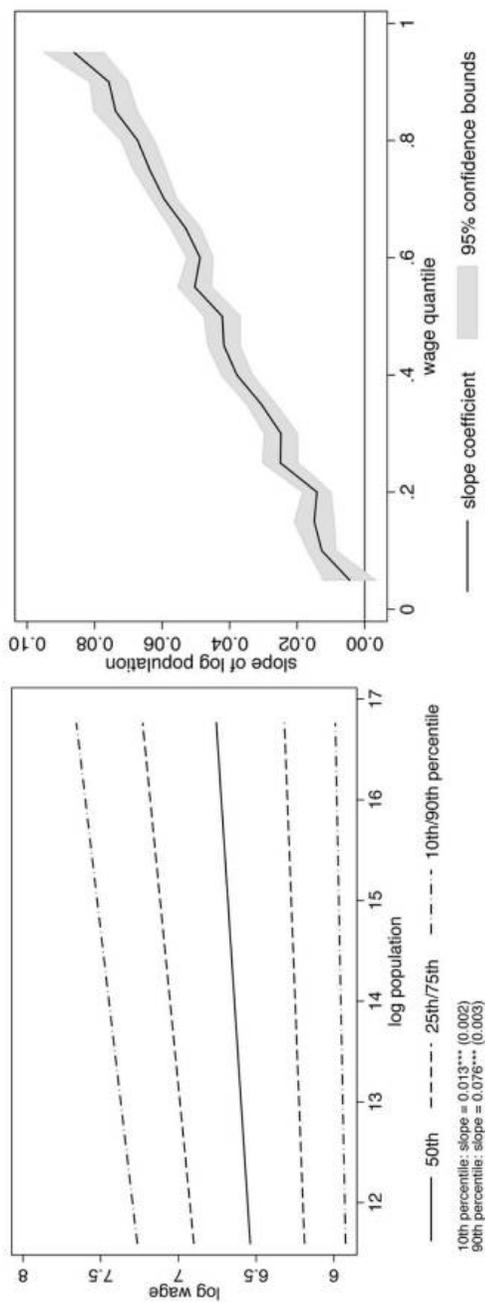


FIG. 2.—Quantile regression of wage on population: *left*, five selected quantiles; *right*, estimated slope for all quantiles

in a larger city is virtually completely absorbed by a 3.8 percent disutility increase due to the cost of living.

The right panel of figure 1 in the introduction shows the entire distribution of skills for full-time earners both in cities with a population of more than 2.5 million and in cities with a population between 100,000 and 1 million. In contrast to the wage distribution, the skill distribution in large cities is only marginally shifted to the right. However, both the upper and the lower tails of the distribution are thicker in the large cities, thus confirming the consistency with the theoretical prediction of thick tails from extreme-skill complementarity.⁹ An explicit look at the tails of the two distributions confirms the thick tail prediction in a statistical sense. The 90th percentile for large cities is 6.99 compared to 6.86 for small cities ($\Delta = 0.132$, $SE = 0.009$, $p < .01$). The 10th percentile for large cities is 5.36 compared to 5.44 for small cities ($\Delta = -0.074$, $SE = 0.006$, $p < .01$). So the large cities have both a significantly lower 10th percentile and a significantly larger 90th percentile, which implies the thick tails.

As with the wage distribution, one could argue that our partition of cities into small and large ones is arbitrary. We therefore also run quantile regressions of our implicit skill measure on city population. Figure 3 visualizes the results of these regressions. It shows that the median (50th percentile) barely changes with city size while the lower percentiles significantly decrease and the upper percentiles significantly increase. This reiterates our finding that the average of skills does not change systematically with city size while the variance of skills increases significantly. The quantile regressions also perfectly account for the top-coding in the wage data up to about the 95th percentile.

VI. Additional Sources of Heterogeneity

In this section we allow for additional heterogeneity in locations, in individual preferences, and in prices. We show that our main theoretical results and empirical findings are robust to allowing for heterogeneity in attractiveness of locations within cities (Sec. VI.A), for nonhomothetic household preferences (Sec. VI.B), and for local price variation of other consumer goods beyond housing (Sec. VI.C).

A. *Heterogeneous Attractiveness of Locations within Cities*

In our analysis, the endogenous choice of housing is a central component. High-skill workers in the same city consume more housing h than

⁹ Note again that the “bumps” in the top tail are due to top-coding; see n. 7. Top codes appear more to the left for large cities because real wages are deflated with higher housing prices.

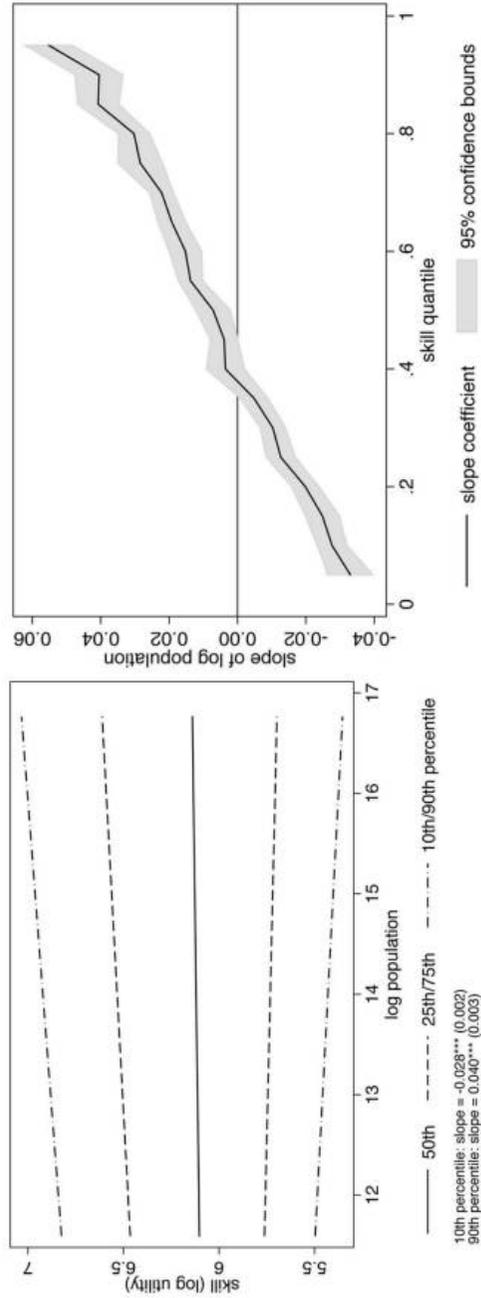


FIG. 3.—Quantile regression of skills (utility) on population: *left*, five selected quantiles; *right*, estimated slope for all quantiles

low-skill workers, and at the same time there is substitution between housing and other consumption goods: same-skilled workers consume less housing in large, expensive cities than in small ones. In order to obtain our wage-based measure of skills, and guided by the theory, we adjust wages by a citywide housing price index, which measures the unit cost (say, per square foot) of housing. To adjust for different choices of quantities, we have obtained that unit cost by means of a hedonic regression that conditions on observables, such as the number of rooms, bathrooms, and so forth. Implicit in this specification is the assumption that all neighborhoods are equally attractive, and citizens with different incomes will therefore share the same neighborhoods, albeit in houses of different sizes.

In reality, though, not all locations within a city are equally attractive. The desirability of neighborhoods depends on factors such as closeness to work, cultural events, restaurants, shopping outlets, and recreational opportunities, as well as on access to good schools and other public goods, or simply on its socioeconomic status. In spatial equilibrium, more attractive locations will have higher housing prices. In the absence of within-city sorting—which will be discussed further down—all households are indifferent between all locations within the city in spatial equilibrium. A simple theoretical framework of this mechanism is the monocentric city model in which the attractiveness of locations within the city decreases with distance from the central business district because of commuting costs.¹⁰ Because there is a trade-off, the housing price now also reflects the attractiveness. The less attractive, the lower the housing price, even for identical agents. As a result, there is a so-called bid-rent function that increases with attractiveness. Only in the city center without any commuting does the price reflect the true cost of living. Anywhere else, the housing price is too low since it embodies both the cost of living and the disutility from less local attractiveness.

We operationalize neighborhoods within cities in the data as public use microdata areas (PUMAs), relatively small areas of around 100,000 inhabitants. We estimate hedonic price indices for all PUMAs across the United States. We then take the average price index of the top 10 percent PUMAs per city, that is, CBSA. The imputed skill distribution based on this maximal housing price index is reported in figure 4. Consistent with our earlier results, the thick tails continue to exist. We tend to see a somewhat bigger tail at the bottom than at the top. Even when we adjust for differential attractiveness of locations within the city, the thick tails result continues to hold.

So far, the logic with differential attractiveness of locations is for identical agents. In fact, housing prices adjust to equalize the difference in

¹⁰ For an overview of the different variants of the monocentric city model and a full characterization, see Fujita (1989).

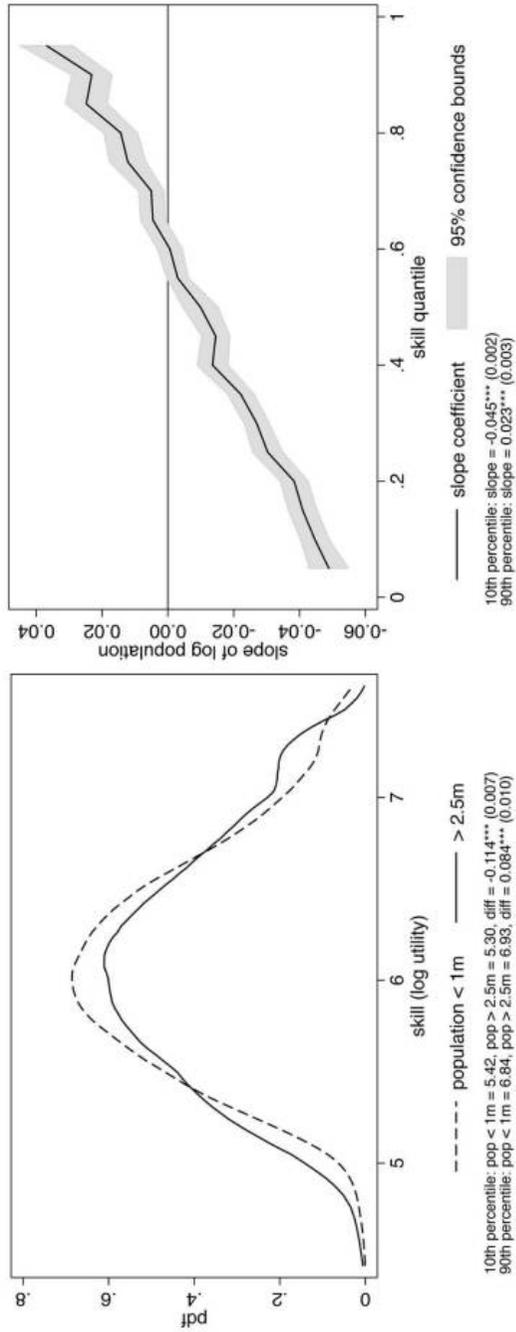


FIG. 4.—Skills based on the average housing price index in the top 10 percent neighborhoods (PUMA) of a city; wage data from the 2009 CPS; *left*, skill distribution for small and large cities; *right*, slopes in quantile regressions of log utility on log population.

attractiveness, and all agents are indifferent where to locate within the city. When, in addition, agents are heterogeneous—as in our model in incomes—and different-income agents value the trade-off between local attractiveness and housing price differently, then there will be sorting within the city. If the opportunity cost of attractiveness is complementary with income, then the richest citizens will sort into the most attractive neighborhoods (e.g., Manhattan in New York or Bloomfield Hills in Detroit). Now there is a bid-rent function for each household type that is observed only where a citizen type actually lives: at the most attractive location for the highest-income types, in intermediate attractive locations for the middle-income types, and at the least attractive locations for the lowest-income types. In the case of perfect sorting, each type would live in a dedicated neighborhood between intersections of the bid-rent functions. The relevant price for utility comparisons across cities would be the bid-rent at the most attractive location, but that is unobserved except for the highest-income types in the most attractive locations. The observed local housing price for less attractive neighborhoods is then a lower bound of the relevant price since the observed price incorporates the cost of commuting.

We estimate this lower bound as the hedonic price index of the PUMA where the observed worker lives. Unfortunately, the wage data from the CPS do not identify the PUMA of the worker. We therefore use wages from the 2009 ACS for this analysis, that is, wage data from the same source as the price data. We then assign each worker in the ACS the housing price of the PUMA where he lives. In figure 5, we first reproduce the basic findings from figure 1. As with CPS data, the ACS data show first-order stochastic dominance in wages and thick tails in skills, though the effect on the lower tail is less pronounced.¹¹

In figure 6, we report the distribution of skills based on the price of the PUMA where the worker lives. Quite remarkably, even in the presence of this biased price index, the thick tails continue to exist. It is not surprising that the lower tail difference is thinner, given that our measure is biased downward, but it is still significant. More importantly, because this housing price index is a lower bound, the actual tail difference must be thicker.

We interpret the exercise in this section as one in which we put bounds on the tails. The neighborhood (PUMA) price index is the lower bound and shifts the distribution too little to the left, resulting in a small lower-tail effect. The maximum price index is the upper bound and shifts the distribution too much to the left, generating a big lower-tail effect but hardly any upper-tail effect. The relevant price index, and hence the distributions, is somewhere in between.

¹¹ The CPS is generally considered the more reliable data source for wage data as the survey is performed personally by phone while the ACS questionnaires are mailed. We therefore use our initial results in fig. 1 as the baseline.

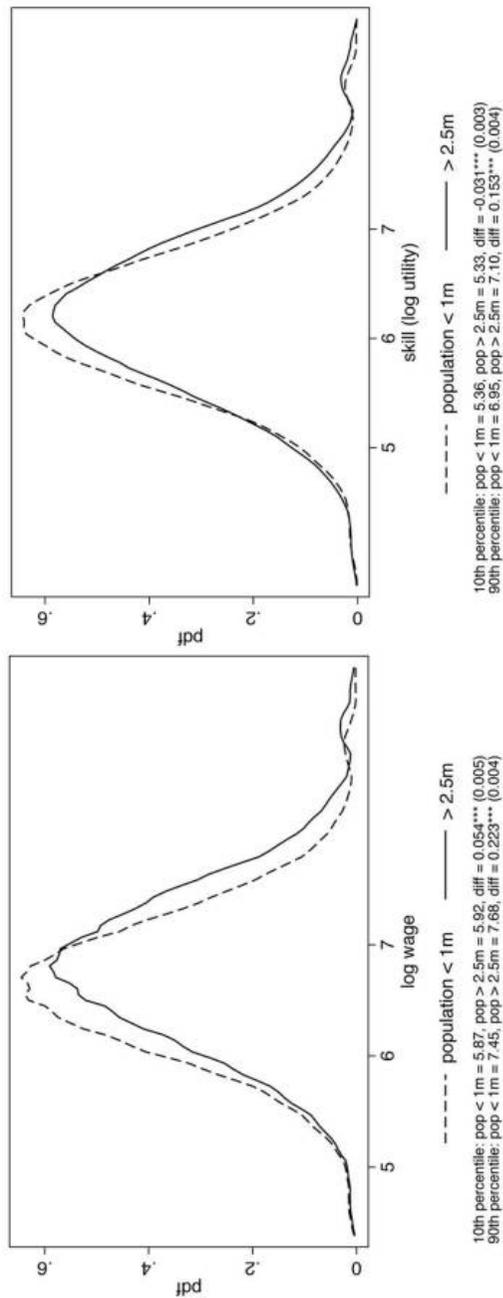


FIG. 5.— Wage and skill distribution for small and large cities; wage data from the 2009 ACS; kernel density estimates (Epanechnikov kernel, bandwidth = 0.1); *left*, wages; *right*, skills.

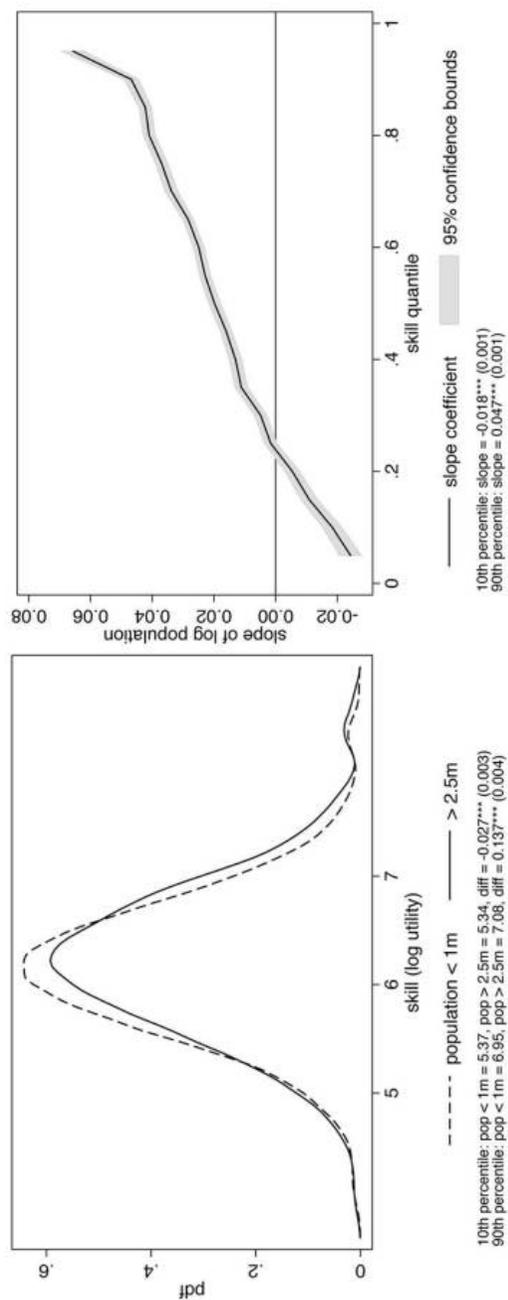


FIG. 6.—Skills based on the housing price in the neighborhood of residence; wage data from the 2009 ACS: *left*, skill distribution for small and large cities; *right*, slopes in quantile regressions of log utility on log population.

We conclude that when we allow for differential attractiveness and ensuing within-city sorting, the thick tail result continues to hold. That does not mean that there is no within-city sorting, quite to the contrary. Figure 7 shows the average housing prices and average log wages across neighborhoods (PUMAs) for five metro areas. The fact that there is a strong correlation between wages and housing prices is indicative of such within-city sorting.

A closer look at the data reveals that the price difference across different neighborhoods within a city is smaller than the price variation across cities: the standard deviation of the PUMA housing price index is 0.36 across and within all CBSAs; the standard deviation *between* CBSAs is 0.33 while it is only 0.25 *within* CBSAs. This is illustrated in figure 8 with the price variation across PUMAs in two cities, New York and Detroit. Among the most expensive residential areas in New York, for example, are the Upper East Side in Manhattan and East Meadow on Long Island. The least expensive areas are East Harlem in Manhattan and the Bronx. Yet, few of the cheapest areas in New York are less expensive (less dark on the map) than the most expensive areas in Detroit. The important point to take away from this is that the housing price differences across cities are relatively large.

B. *Nonhomothetic Household Preferences*

So far we have assumed homothetic (Cobb-Douglas) preferences over housing and consumption. Motivated by the empirical finding of Davis and Ortalo-Magné (2011) that housing expenditure is, on average, remarkably constant across different cities, we have used their estimated expenditure share on housing of $\hat{\alpha} = 0.24$. Yet, even if the average expenditure share of housing is constant across cities of different size, there may well be variation across individuals of different incomes. As a result, the Engel curve that relates expenditure to income is no longer linear as it is under Cobb-Douglas. Below we show that there is indeed evidence in our data of a concave Engel curve: the rich spend proportionally less on housing.

Nonhomothetic preferences have important consequences for both our theoretical model and our empirical strategy. First, decreasing housing expenditure shares with respect to income introduce an alternative mechanism for sorting across cities as high-skill workers care less about local housing prices than low-skill workers.¹² Second, our price-based skill measure derived from the homothetic Cobb-Douglas preferences needs to be adjusted.

¹² For example, Schmidheiny (2006) studies within-city sorting from assuming non-homothetic (Stone-Geary) preferences.

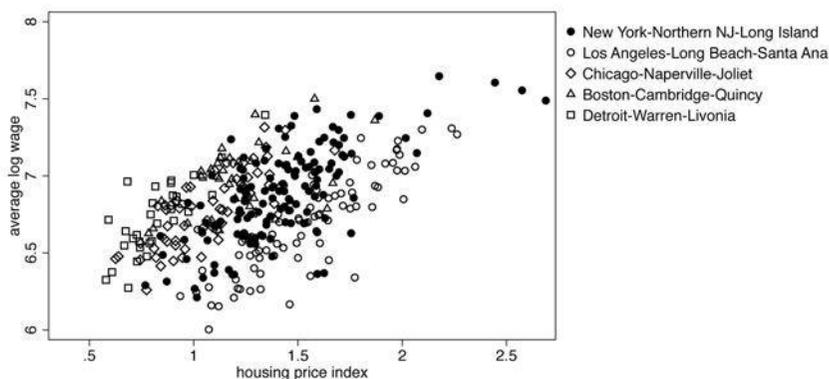


FIG. 7.—Average rental housing prices and average log wages across PUMAs in five selected CBSAs.

It is straightforward to introduce nonhomothetic preferences into our theoretical framework. We follow the most common way in the literature and model it by means of Stone-Geary preferences. They can be written as $u(c, h) = c^{1-\alpha}(h - \underline{h})^\alpha$, where \underline{h} is the subsistence level of housing, and housing consumption is restricted to $h \geq \underline{h}$. Given housing prices p and the budget constraint $c + ph \leq w$, from the first-order conditions, optimal expenditures on housing and consumption can be written as $ph^* = \alpha w + (1 - \alpha)p\underline{h}$ and $c^* = (1 - \alpha)(w - p\underline{h})$, with the indirect utility given by

$$u(c^*, h^*) = (1 - \alpha)^{1-\alpha} \alpha^\alpha p^{1-\alpha} \left(\frac{w}{p} - \underline{h} \right). \quad (16)$$

Assuming the CES production technology and Stone-Geary preferences with $\underline{h} > 0$, our model predicts first-order stochastic dominance of the skill distribution in large cities in simulations.¹³ Hence, nonhomothetic preferences do not generate thick tails per se. But allowing for extreme-skill complementarities in addition, our model still predicts thick tails, for \underline{h} positive but small. In the next paragraph we therefore explore which mechanism prevails empirically.

When housing expenditure varies by income, the utility and therefore our measure of skill must be adjusted. Assuming Stone-Geary preferences, the expenditure share on housing is a linear function in the inverse of wages:

¹³ Matlab code for the simulations can be obtained on request.



FIG. 8.—Average rental housing prices across PUMAs: *left*, New York–Northern New Jersey–Long Island CBSA; *right*, Detroit–Warren–Livonia CBSA. The housing price index is normalized to one for the US sample average.

$$\frac{ph^*}{w} = \alpha + (1 - \alpha)\underline{h}\frac{p}{w}.$$

This gives us a regression of the housing expenditure share s_i on p_j/w_i :

$$s_i = \alpha + \beta \frac{p_j}{w_i} + \varepsilon_i, \quad (17)$$

where $s_i = p_j h_i^*/w_i$. The parameter α is estimated directly while the parameter \underline{h} is estimated as $\hat{\underline{h}} = \hat{\beta}/(1 - \hat{\alpha})$.

We use individual data on expenditure shares (see App. B for details) from the Consumer Expenditure Survey (CEX) to estimate the two parameters α and \underline{h} . We obtain $\hat{\alpha} = 0.224$ (SE = 0.005) and $\hat{\underline{h}} = 27.7$ (SE = 3.8). The implied expenditure shares vary considerably, from 35 percent for low-income households to 22 percent for high-income households as graphed in the left panel of figure 9. The panel also shows that the functional form assumed by Stone-Geary fits the data astonishingly well. Yet, the varying expenditure share and the resulting nonlinearity of the Engel curve do not substantially alter the emergence of thick tails. The right panel of figure 9 shows the resulting skill distribution, which has properties very similar to those under Cobb-Douglas preferences. If anything, the thick tail differences in both the lower and the upper tails are slightly more pronounced. This shows that our evidence for extreme-skill complementarities still holds even after accounting for nonhomothetic preferences.

C. Variation in Consumption Prices

Local prices are crucial for our strategy to back out skills from observed nominal wages. In this subsection we look not only at local housing prices but also at local prices of consumption goods. It may well be that consumption prices in large cities are systematically higher than in smaller cities, thus adding further to the real cost of living in large cities. We use the ACCRA Cost of Living Index from C2ER. See Appendix B for details. The variation in consumption prices is substantially lower than in housing prices (the standard deviation across metropolitan areas is 30.1 for the housing prices index compared to 9.1 for grocery items, 14.2 for utilities, 5.6 for transport, 8.2 for health, and 4.6 for services; all price indices are normalized to mean 100).

Figure 10 plots the distribution of skills for large and small cities. The measure is wages adjusted for local price differences in all goods categories reported in the ACCRA data, including housing, consumption

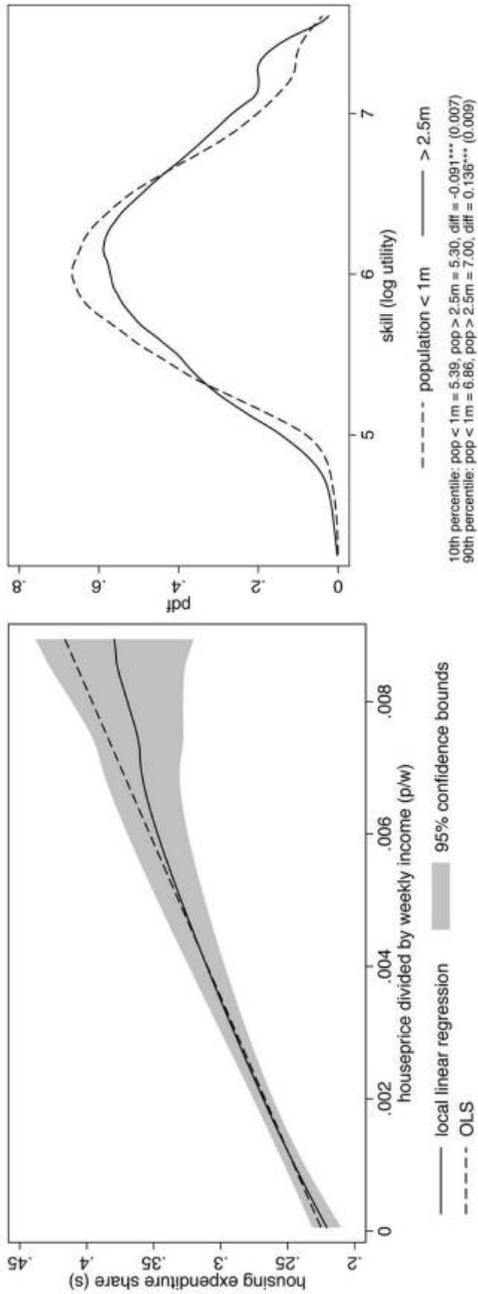


FIG. 9.—Expenditure share (*left*) and skill distribution (*right*) using Stone-Geary preferences, $\hat{\alpha} = 0.224$, $\hat{\lambda} = 27.7$

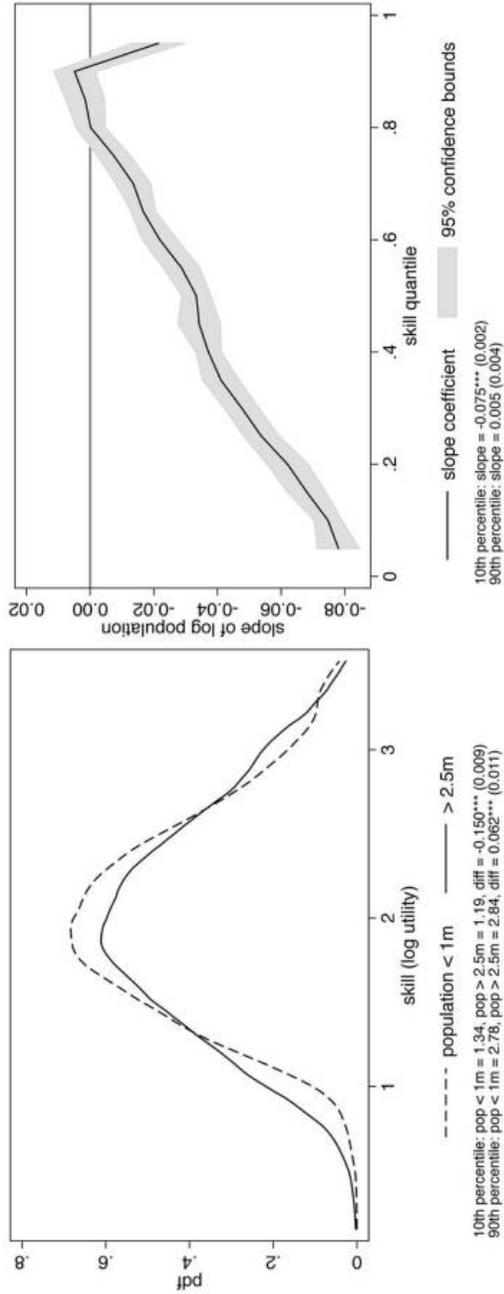


FIG. 10.—Skill distribution using ACCRA Cost of Living Index, adjusting for variation in prices for all goods: *left*, skill distribution for small and large cities; *right*, slopes in quantile regressions of log utility on log population.

goods, and services.¹⁴ When including the price index for all consumption and housing, we find that the left tail difference becomes more pronounced while the right tail difference is less so. This indicates that consumption prices are systematically higher in larger cities, but to a limited extent, since this effect does not annihilate the existence of fat tails. Note again that the third crossing at the very top is an artifact of the top-coding (see n. 7).

These findings should be interpreted with some caution and a few caveats. First, the quality of the ACCRA data is dubious.¹⁵ Second, even within a given location, there could be variation in consumption prices paid by skill level. For example, because of different search intensity, the existence of locally segregated markets, and so forth, the low-skilled may end up paying different prices for similar goods within the same city. Using scanner data on household purchases, Broda, Leibtag, and Weinstein (2009) find that the poor pay less. Third, data consisting of price indexes and price surveys are likely to not fully account for quality and diversity differences. Because of their size, large cities have more variety on offer, and the quality of goods may differ substantially across different cities. Even if a consumer pays higher prices, a price index incorporating the diversity and quality on offer will be lower.¹⁶ In addition, ACCRA puts more weight on housing than we do in our baseline result (see details in n. 14). Using our budget share, the skill distribution of large cities would be

¹⁴ ACCRA reports a composite price index, which is the weighted average of the six subindices, i.e.,

$$P_{\text{composite}} = \alpha_{\text{grocery}} P_{\text{grocery}} + \dots + \alpha_{\text{services}} P_{\text{services}},$$

where the α 's are the expenditure shares of the six categories summing up to one. We do not use this aggregation as it is inconsistent with Cobb-Douglas utility. Instead, we use

$$P_{\text{composite}} = (P_{\text{grocery}})^{\alpha_{\text{grocery}}} \dots (P_{\text{services}})^{\alpha_{\text{services}}}.$$

The implied skills are calculated as

$$U_i = w_j / (P_{\text{grocery},j}^{0.1249} \cdot P_{\text{housing},j}^{0.2918} \cdot P_{\text{utilities},j}^{0.0998} \cdot P_{\text{transport},j}^{0.111} \cdot P_{\text{health},j}^{0.0406} \cdot P_{\text{services},j}^{0.3319}).$$

The exponents are budget shares taken from the ACCRA Cost of Living Index Manual, November 2009 version (current versions of the manual are available online at <http://www.coli.org/surveyforms/colimanual.pdf>). The ACCRA values differ from the ones we use for the baseline results of this paper. On the basis of Davis and Ortalo-Magné (2011), we use a budget share of 0.24 for housing rather than the 0.29 used by ACCRA.

¹⁵ Koo, Phillips, and Sigalla (2000) discuss several problems of the ACCRA data.

¹⁶ This also appears to be an issue when studying price differences across different countries. Comparing the results of price differences across borders, Broda and Weinstein (2008) find that significant price differences that are found using price indexes are not replicated once they use US and Canadian barcode data. Their work is supportive of simple pricing models in which the degree of market segmentation across the border is similar to that within borders.

shifted less to the left, and the lower-tail difference would be less pronounced and hence more in line with our baseline results. We therefore see the results in figure 10 as a very conservative upper bound of how the inclusion of consumption price differentials affects our initial findings.

VII. Decomposing Observable and Unobservable Skills

Our measure of skills is a price-based measure, calculated as the residual of wages after adjusting for housing prices. As such, it is a comprehensive measure that encompasses both observed and unobserved characteristics. In this section we verify the robustness of the spatial sorting result using observed, direct measures of skill. Not only does it allow us to decompose the contribution of observed and unobserved characteristics, but it also makes explicit if and why the results with our comprehensive measure are different from a long tradition of previous literature. We show that our thick tail result also holds using directly observed measures of skill. We also contribute to the debate on the role of noncognitive skills and how much they contribute to wage determination. Here, we focus explicitly on the impact of noncognitive skills across labor markets of different size.

In the first instance (Sec. VII.A), we use years of schooling, occupation, and industry as direct measures of skills instead of our wage-based measure. We then investigate the role of mobility and migration that is systematic by nationality (Sec. VII.B) and also whether there are any location decisions determined by age over the life cycle (Sec. VII.C). Finally, we quantify the impact of those observables simultaneously (Sec. VII.D) and decompose the tails into an explained and an unexplained component.

A. *Observed Measures of Skills: Education, Occupation, and Industry*

As a first robustness check and as external validation, we compare our implicit skill distribution with that of educational attainment. The top-left panel in figure 11 shows the distribution of highest educational attainment for the same CPS population as our wage data, where workers are grouped in five education categories. The same pattern as with our implicit measure arises: both the highest- and the lowest-skilled workers are disproportionately more frequent in larger cities (population above 2.5 million) than in smaller ones (population between 100,000 and 1 million). What is most striking about this observation is that the thick tails in the distribution of educational attainment are obtained independently of how we constructed our measure of skills before. Here, no theory is needed and the measure of skills is determined exogenously. Using ob-

servable, self-reported measures of skills, we find a distribution with thicker tails in larger cities, both in the aggregate and at the individual city level.

In principle, individual skills can be decomposed into an observed component (e.g., education) and an unobserved component (e.g., ability). We already know that the observed component indeed exhibits thick tails. To get to the unobserved component, we regress our implicit skill measure (log utility) on dummy variables for all 16 observed education cate-

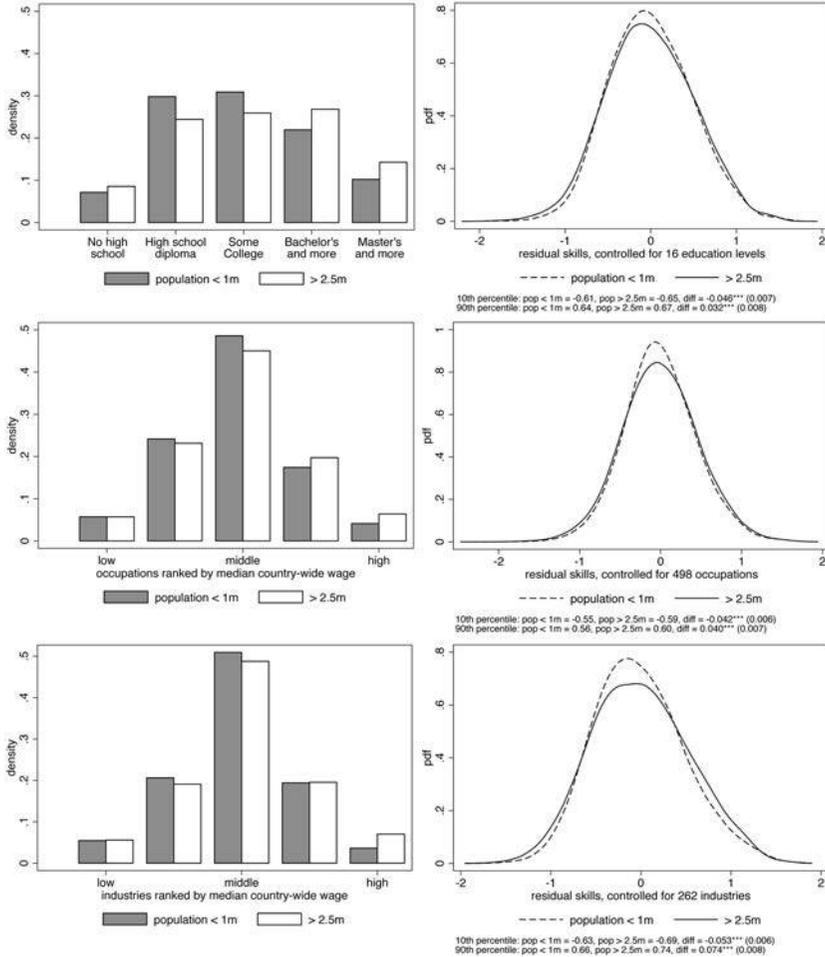


FIG. 11.—Observed and residual skills of full-time wage earners in the 2009 CPS for small and large cities: *top*, education, highest completed grade; *middle*, occupations, three-digit 2000 Standard Occupational Classification code; *bottom*, industry, three-digit North American Industry Classification System code.

gories.¹⁷ The residual of this regression is the “residual skill” after controlling for observed education. A high value means that the worker is more skilled relative to other workers with the same education. This can be a very successful lawyer or a high school dropout who becomes a successful entrepreneur.

The top-right panel in figure 11 shows the resulting distribution of residual skills for large and small cities. We see that the thick tails persist for the unobserved component and are more pronounced at the bottom than at the top. Large cities seem to host more of the relatively low-skilled even after controlling for education category. Taken together, these two pieces of evidence mean that large cities attract relatively more of the best-educated and more of the least-educated workers. In addition, they attract the more talented, given their education level. We also see that there is an asymmetry between the upper and the lower tails. This indicates that the sorting of the most talented workers is based more on the observed component of skills (education) while the sorting of the least talented is based more on the unobserved component, a finding that we will reiterate in the next paragraph, where we condition on occupations.

Like education, occupation can be interpreted as a direct measure of skills. More highly skilled individuals are likely to be employed in higher-ranked occupations. If we can find a way to rank occupations, it can give us some insights into the sorting of workers based on this observed characteristic. The CPS contains the occupation of workers, in addition to wages and education. This occupational classification has been used as an alter-

¹⁷ We estimate censored (Tobit) regression, accounting for the top-coding of the wage data. We regress $\log(u_{ij})$ on a constant and a set of dummy variables for education with basic education as the reference group. This dummy variable regression is fully consistent with our theoretical model. Recall that the wage ratio of skill type i relative to skill type 1 is constant across cities,

$$\log(w_{ij}/w_{1j}) = \log(w_{ij}) - \log(w_{1j}) = \beta_i,$$

and therefore the ratio of log utility, too,

$$\log(u_{ij}/u_{1j}) = \log[(w_{ij}/p_j^\alpha)/(w_{1j}/p_j^\alpha)] = \log(w_{ij}/w_{1j}) = \beta_i.$$

The log utility of skill type i can therefore be expressed as

$$\log(u_{ij}) = \beta_1 + \beta_2 \times d_2 + \dots + \beta_i \times d_i + \dots,$$

where β_1 is $\log(u_{ij}) = \log(u_{i1}) = \log(u_i)$ of the reference skill type 1, which is constant across cities j . Notice that regressing $\log(\text{wage})$ on a constant and dummy variables for education would not be consistent with our theoretical model as the constant, $\log(w_{1j})$, would be city specific. The usual wage regression therefore needs city fixed effects, which we do not need.

native direct measure of skills before (see, e.g., Autor, Levy, and Murnane 2003). The CPS reports occupations in 498 categories from the 2000 Standard Occupational Classification by the US Bureau of Labor Statistics. We use countrywide median wages by occupation category as a proxy of their skill requirement (this approach is used in, e.g., Goos, Manning, and Salomons [2009]). We then group the 498 occupation categories into five groups: group 1 includes the lowest-paid occupations covering 5 percent of all workers; group 2, the next 20 percent; the middle group 3, 50 percent of workers; group 4, the next 20 percent; and the high-skill group 5, the highest-paid occupations covering 5 percent of all workers. The low-skill group 1 includes occupations such as dishwashers, waiters, and child care workers; the middle group 3 includes occupations such as secretaries and truck drivers; and the high-skill group includes chief executives, surgeons and lawyers. The middle-left panel in figure 11 shows the distribution of these five occupation groups separately for small cities and for large cities. It shows that workers in the highest-paid occupations locate relatively more often to large cities while the middle occupations locate more to small cities. The lowest-paid occupations are equally frequent in small and large cities. Note that already the fact that the lowest group is not more frequent in small cities is a contradiction of the first-order stochastic dominance observed in nominal wages. This direct evidence of thick tails is very similar to the result we found for education groups. The effect on the lower tail is less pronounced than for the upper tail.

As with the analysis of education, we next decompose skills into a component observed through occupations and the residual unobserved component. We regress our skill measure (log utility) on dummy variables for all 498 occupation categories.¹⁸ The residual of this regression is the “residual skill” after controlling for observed occupation. The middle-right panel in figure 11 shows the resulting distribution of residual skills for large and small cities. We see that the thick tails persist for the unobserved component. In particular, large cities seem to host more of the relatively low-skilled in each occupation than small cities. As a result, large cities attract relatively more workers in the highest-paid occupations and more in the lowest-paid occupations. In addition, they attract the more talented, given their occupation. We also see that the sorting of the most talented workers is based more on the observed component of skills (occupation) while the sorting of the least talented is based more on the unobserved component, a finding that we already made using education as a direct measure.

¹⁸ We estimate a censored (Tobit) regression accounting for the top-coding of the wage data.

A potential threat to our identification of skills is the different industrial composition of cities. Wages systematically differ across industries and the mobility of workers varies across sectors (see, e.g., Davis, Faberman, and Haltiwanger 2006). At the same time, industrial composition changes across cities: some cities specialize in particular industries and other cities have a diverse industrial mix. If industry composition varies systematically with city size, this could be an alternative explanation for our finding of thick tails. We therefore seek to control for the wage component related to the industry. As in the analysis of occupations, we first rank the 262 industries by countrywide median wages and group them into five groups: group 1 includes the lowest-paid industries covering 5 percent of all workers; group 2, the next 20 percent; the middle group 3, 50 percent of workers; group 4, the next 20 percent; and the high-skill group 5, the highest-paid industries covering 5 percent of all workers. The middle-left panel in figure 11 shows the distribution of these five industry groups separately for small cities and for large cities. We see indeed that large cities attract more workers from industries that pay the highest wages as well as more workers from industries that pay the lowest wages.

As we did with education and occupation, we decompose our implicit skill measure into a component observed through industries and the residual unobserved component. We regress our skill measure (log utility) on dummy variables for all 262 four-digit industries.¹⁹ The residual of this regression is the “residual skill” after controlling for the industry the worker operates in. The bottom-right panel in figure 11 shows the resulting distribution of residual skills for large and small cities. We see substantial thick tails on both the upper and the lower ends when controlling for industries. This shows that while the industrial composition may vary across cities, it does not do so systematically across small and large cities.

B. Migration

Casual observation suggests that large cities tend to have a disproportionate representation of low-skilled immigrant workers. Often kitchen staff in restaurants or construction workers are immigrants with low skills and incomes. And indeed, while the foreign-born are overall a relatively small fraction of the working population (less than 10 percent), the data confirm that they are much more likely to locate in large cities (12 percent of the work force) than in small cities (5 percent). Maybe the effect

¹⁹ We estimate a censored (Tobit) regression accounting for the top-coding of the wage data.

of disproportionate representation of the low-skilled in large cities is driven by immigration.

In the context of our model it does not matter whether it is low-skilled Americans or low-skilled immigrants who disproportionately locate to large cities. In equilibrium they should be indifferent. Of course, there is likely to be within-skill heterogeneity (in preferences, e.g.), and some low-skilled workers will strictly prefer to locate to either large or small cities. Thus it may well be the case that migrants have certain benefits from locating to large cities. For example, networks (see Munshi 2003) play an important role in the location decision of migrants, and if only migrants have that benefit, at a competitively set wage, migrants will strictly prefer to locate in the city that offers the same utility plus the network benefit. Alternatively, migrants may locate in large cities because of limited information about smaller cities.

Consider this logic with preference bias in the context of our baseline model. There are three skill types (low, middle, and high) and two cities (small and large). Assume now that all migrants are low-skilled and have a preference for the large city; that is, they get additional utility from living there. Suppose that in an equilibrium without migrant preferences, the number of low-skilled workers in the large city is m_{11} . Now with the migrants' preference for large cities, all migrants will want to live in the large city. As long as the number of migrants is below m_{11} , the general equilibrium will not change at all. All low-skilled migrants will live in the large city, some low-skilled natives will also live in the large city, and all the others will live in the small city. The native workers remain the marginal workers and are still indifferent between the large and the small city. Migrants, however, strictly prefer the large city and are not indifferent. The resulting equilibrium skill distribution is not affected by migrants. Just looking at migrants, we would see a (very much) thicker lower tail in large cities. Just looking at natives, we would see a less thick (if not thinner) lower tail in big cities. The important implication is that even if we observed those location biases, it would not affect the aggregate predictions.

In general, as long as the migrants—or any subgroup of the population for that matter—of a particular skill and with a taste for a particular city do not outnumber the equilibrium number of workers of that skill in that city, our model predictions are unchanged by migrant preferences for large cities. The fact that we observe a substantial number of natives of all skill types (including the very low-skilled) in all cities provides empirical support that the native workers are the indifferent marginal workers for all skill types.

Empirically, we even find no evidence of a systematic bias in the location decision of the migrants. To evaluate this, we split the sample up

into natives and foreign-born workers. Figure 12 reports the plot of both distributions. Not surprisingly, the implied skill distribution for the foreign-born is more skewed to the left than that of the natives. We also see that the distribution of foreign-born workers has thick tails, for both the low- and the high-skilled. The latter is maybe most surprising: not only do the low-skilled foreign-born disproportionately migrate to large cities, but so do the high-skilled migrants. Most importantly, even after all migrants are subtracted, the distribution of natives has thicker tails in large cities. The thick tails are therefore not driven by selective migration decisions by nonnatives.

C. Age: Location Decisions over the Life Cycle

One plausible mechanism, and different from the technological one we propose, is that the spatial sorting pattern is driven by location decisions that vary over the life cycle. We distinguish between three candidate mechanisms in the presence of human capital accumulation. First, age-dependent preferences could lead to variation in the location decision over the life cycle. Young people prefer the excitement of the city, while older people settle for a quiet life. Second, family and marriage considerations determine whether to live in an urban or rural environment, for example, as in Gautier, Svarer, and Teulings (2009). Singles find a better marriage market in the big city, while those married with children look for green spaces and schools for the children. Third, labor market learning and human capital accumulation may affect the location decision (as in De la Roca and Puga 2012). Young workers try their luck in the big city, starting off at low salaries. At a later age, those who have learned (or are lucky) to be very productive stay and earn high salaries, while the unlucky who have learned little and have added limited human capital return to their small town at moderate salaries.

Each of these three mechanisms induces a systematic spatial sorting pattern and, as a result, a systematic skill distribution for the entire cross section. This can lead to thick tails or first-order stochastic dominance. Most importantly, if it is life cycle driven, it will differ for different age groups. For example, there may be stochastic dominance of the small cities for the young and stochastic dominance of the large cities for the old, thus leading to thick tails overall. For that purpose, we split the sample into four different age cohorts and investigate the tail properties of the skill distribution for each cohort as reported in figure 13. For each of the cohorts, there are thick tails, and the statistics show that they are highly significant. While we obviously cannot rule out that any of the mechanisms mentioned above is at work, this evidence lends support to the fact that these mechanisms are not sufficiently strong to undo the impact of the extreme-skill complementarities.

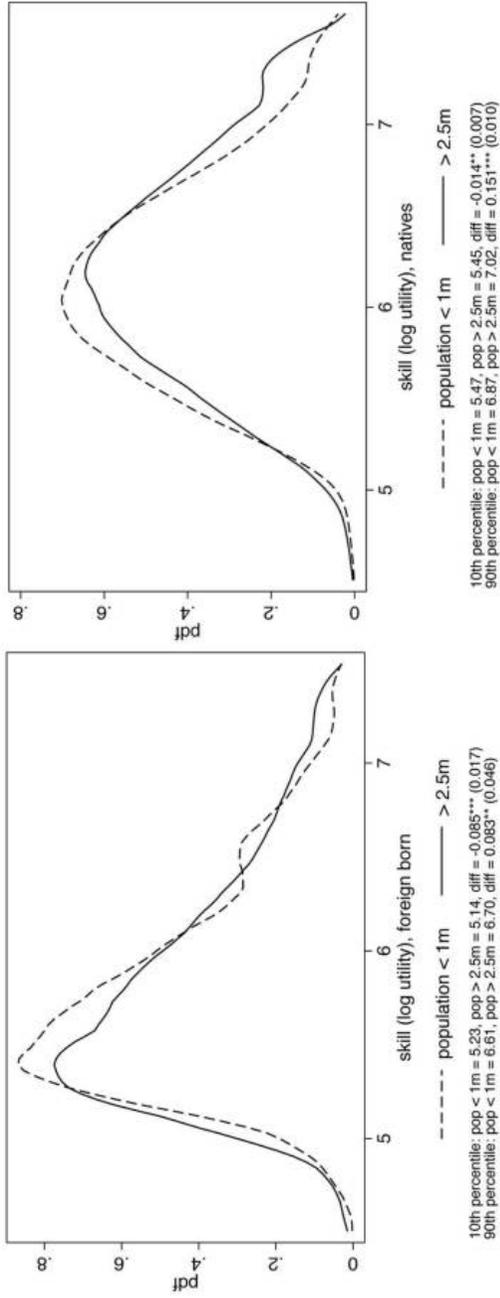


FIG. 12.—Skill distribution for small and large cities: *left*, foreign-born workers; *right*, natives

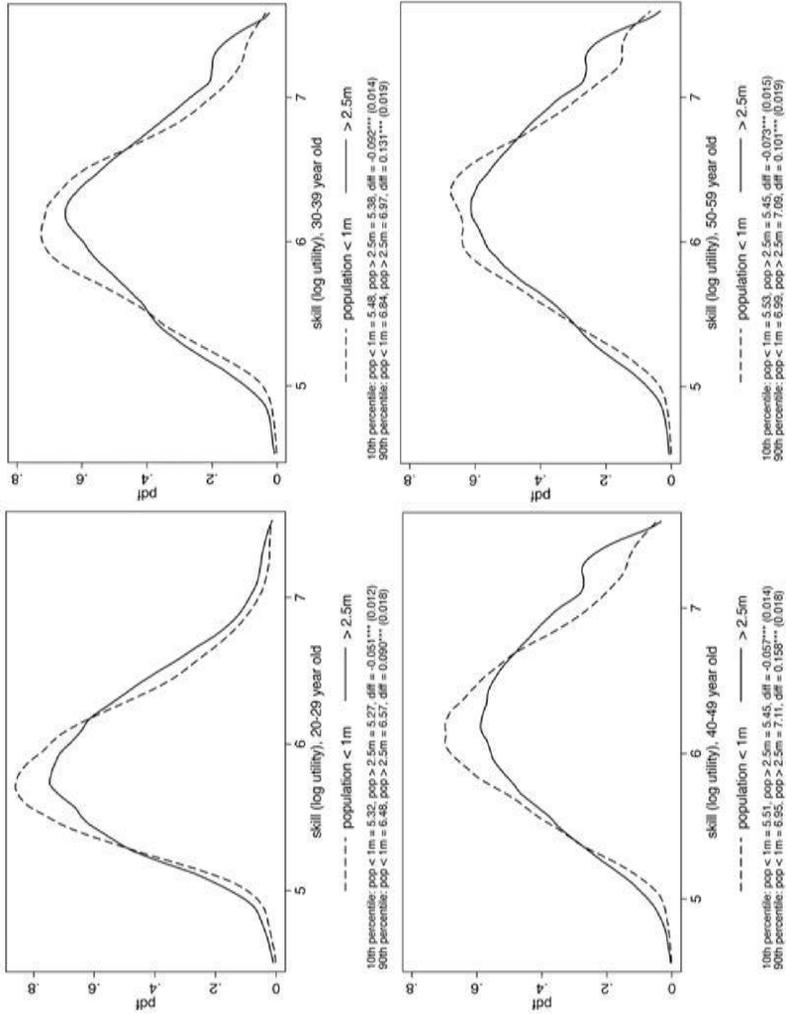


FIG. 13.—Skill distribution for small and large cities by age group

D. Decomposition by Education, Occupation, Industry, and Individual Attributes

This subsection seeks to explain the differences in the tails of the skill distribution simultaneously by all three variables in Section VII.A (education, occupation, industry) plus individual attributes (sex, age, race, foreign birth). In regression models explaining the mean, this can easily be done by the Oaxaca-Blinder decomposition (Oaxaca 1973; Blinder 1973). However, we are interested in explaining differences in the tails of the skill distribution, which is a much harder task. Fortunately, there is a fast-evolving econometrics literature on the decomposition of entire (wage) distributions (Juhn, Murphy, and Pierce 1993; DiNardo, Fortin, and Lemieux 1996; see Fortin, Lemieux, and Firpo [2011] for a review).

We are using two very recent approaches. The first approach is based on Chernozhukov, Fernández-Val, and Melly (2013). This approach estimates the entire distribution of skills conditional on the observed covariates for both groups (small and large cities) separately using quantile regressions. They then integrate the conditional distributions over the covariates to get the predicted marginal distribution of each group. With this they can predict counterfactual marginal distributions such as the distribution of skills in the large city, given that it had the same distribution of covariates as the small city. We refer to the difference between the marginal distribution and the counterfactual distribution as the “unexplained” difference.²⁰

We summarize the results by reporting the impact on the 10th and 90th percentiles. Chernozhukov et al. (2013) cannot easily decompose the explained difference into different sets of covariates. We therefore also apply an alternative decomposition proposed by Firpo et al. (2009). Their approach, based on so-called rescaled influence functions, allows them to conveniently decompose the explained differences in the quantiles into the contribution of each covariate. A downside of their study is that the basic approach is less intuitive.²¹

The original classification into 16 education categories, 498 occupations, and 262 four-digit industries would lead to over 900 parameters

²⁰ In labor economics, the unexplained difference is often called a “wage structure effect.” The reason is that the difference in the conditional distributions between the two groups may stem from different wage schemes, i.e., returns to, e.g., education. However, it may also stem from different conditional (i.e., residual) skill distributions in the two groups. The decompositions by Firpo, Fortin, and Lemieux (2009) and Chernozhukov et al. (2013) do not and cannot disentangle wage structure from residual skills. We do not take a stance at either interpretation and simply call it the unexplained difference.

²¹ The method of Firpo et al. (2009) is an approximation for a marginal location shift of the distribution of the covariate. It is not known how good this approximation is if the change in the distribution of the covariate of interest is large or if the covariate of interest is discrete (like dummy variables). See Rothe (2012, app. B) for a discussion.

in each city group, which are hard to identify with our data. We therefore use two-digit industry classifications (52 classes) and two-digit occupation codes (22 categories), which are both assigned by the National Bureau of Economic Research. Age enters as a second-order polynomial and race in four groups (white, black, Asian, and other).

Table 1 shows the results of the two approaches. The first three rows report the raw sample quantiles of the skill, $\log(\text{utility})$, distributions as in the right panel of figure 1. Both methods predict these quantiles well. With the method of Firpo et al. (2009), only about 26 percent of the predicted difference in the 10th percentile can be explained by the composition of observed characteristics. However, 46 percent of the predicted difference in the 90th percentile can be explained by observed characteristics. Most of this is explained by education (36 percent) and occupation (17 percent). The strong explanatory power of observables for the top tail and the relatively low explanatory power for the bottom tail reiterate the findings in the previous sections. The method of Chernozhukov et al. (2013) leads to very similar results: 57 percent of the 90th percentile is explained by the composition of observables and 28 percent of the 10th percentile.

The novel finding here is that there is an asymmetry between the low- and the high-skilled. For the low-skilled, very little of the difference between big and small cities can be explained by observables, whereas for the high-skilled, about half can be explained by observables.

VIII. Quantifying the Production Technology

In this section, we use the equilibrium properties of the theoretical model to quantify the features of the technology. The model allows us to obtain quantitative information on the technology, in particular the differences in TFP across big and small cities, the productivity differences across differentially skilled workers, as well as the magnitude of the extreme-skill complementarity.

Using the observed skill distributions, we perform a very simple quantitative exercise to get an idea of the magnitude of these underlying parameter values of the production technology. To that end, we can actually solve the system of equations explicitly (see App. A for the derivation) to obtain a system of five equations in five unknowns: λ , A_1 , A_2 , y_1 , and y_3 , where y_2 is normalized to one and γ is exogenously chosen.²²

To adapt the data to our model with three skill types and two cities, we partition the distribution of our wage-based measure of skills into three

²² In a CES technology, there is an indeterminacy between the A_j 's and the y_j 's. Here, since the skill $i = 2$ is CES with the composite of skills 1 and 3, it is proportional to A_j (see the second equation), and we can normalize it to one.

TABLE 1
DECOMPOSING THE SKILL DISTRIBUTIONS OF LARGE AND SMALL CITIES

	10% Quantile	90% Quantile
Observed quantiles:		
Large cities	5.365*** (.004)	6.994*** (.006)
Small cities	5.439*** (.005)	6.862*** (.007)
Difference	-.074*** (.006)	.132*** (.009)
	Firpo et al. (2009)	
Predicted quantiles:		
Large cities	5.387*** (.005)	7.022*** (.005)
Small cities	5.454*** (.004)	6.878*** (.008)
Difference	-.068*** (.007)	.144*** (.009)
Explained by observables:		
Education (16 categories)	.003** (.002)	.052*** (.002)
Occupation (22 categories)	.004* (.002)	.025*** (.003)
Industry (51 categories)	-.001 (.002)	.013*** (.002)
Race (4 groups)	-.004*** (.001)	-.015*** (.001)
Sex	-.001* (.001)	-.002* (.001)
Foreign-born	-.020*** (.002)	-.004*** (.001)
Age (2nd-order polynomial)	.000 (.001)	-.002* (.001)
Total explained by observables	-.018*** (.004)	.067*** (.005)
Not explained by observables	-.049*** (.006)	.077*** (.008)
	Chernozhukov et al. (2013)	
Predicted quantile difference	-.068*** (.006)	.113*** (.009)
Explained by observables	-.019*** (.004)	.064*** (.005)
Not explained by observables	-.050*** (.007)	.049*** (.007)

SOURCE.—2009 CPS data on 25,584 workers in 202 small CBSAs and 34,999 workers in 21 large CBSAs.

NOTE.—Large cities: population > 2.5 million; small cities: population < 1 million. Bootstrapped standard errors are in parentheses.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

types $i = 1, 2, 3$ corresponding to the 20-60-20 percentiles and construct two city types $j = 1, 2$, type 1 with population larger than 2.5 million and city type 2 with population less than 1 million. We then use observable median wages w_{ij} within each city type and skill group, the actual number of agents of each type in each city m_{ij} , and the number of cities in the sample C_j .²³ Housing prices p_j satisfy equation (4).²⁴ We use the param-

²³ We use median wages because the tails are truncated lognormals and thus are heavily skewed. We adjust the observed wages such that the relative wages of different skill types are constant across cities as the theory predicts. Also, we impose symmetry on the observed numbers for skill types 1 and 3.

²⁴ The observed prices for the two city groups are very similar to that implied by relative wages.

TABLE 2
 QUANTIFYING THE PRODUCTION TECHNOLOGY

OBSERVED MODEL OUTCOMES							
CITY j	w_{1j}	w_{2j}	w_{3j}	m_{1j}	m_{2j}	m_{3j}	C_j
1	416	844	1,923	730,509	1,953,303	730,509	21
2	354	717	1,634	30,900	105,516	30,900	204
IMPLIED PRODUCTION TECHNOLOGY FOR DIFFERENT VALUES OF γ							
γ	λ	A_1	A_2	y_1	y_2	y_3	
.655	1.0407	190,228	59,107	.2329	1	1.0762	
.8	1.0193	19,118	9,065	.3189	1	1.4733	
.9	1.0086	3,992	2,534	.3964	1	1.8317	

NOTE.—The term w_{ij} is the weekly median wage (in US dollars); m_{ij} is the number of workers of skill type i in cities of type j (in units); C_j is the number of cities of type j (in units); γ is the exogenously chosen technology parameter; and λ , A_j , and y_i are the estimated technology parameters.

eter values $\alpha = 0.24$ (the share of expenditure on housing) and estimate the production technology parameters for different values of γ . The results are reported in table 2. As will become clear next, we tend to find more reasonable estimates for values of γ closer to one.

For $\gamma = 0.655$, which corresponds to an elasticity of substitution between low- and high-skill workers of 2.9 (Acemoglu and Autor 2011), TFP is three times higher in the large city than in the small city, while the top skill group is 4.6 times as productive as the bottom skill group, but merely 8 percent more productive than the middle skill types. The parameter λ that measures the extent of the complementarities/substitutabilities is larger than one, confirming that there are complementarities between the extreme skill types. The magnitude of λ is 1.041. Instead, for higher γ , there is less curvature on the amount of labor in output produced, and as a result, TFP differences are much closer. For $\gamma = 0.8$, TFP is double, and the high-skilled are still 4.6 times more productive than the low-skilled (this is also the case for $\gamma = 0.9$), but now they are 47 percent more productive than the middle-skilled. For $\gamma = 0.9$, TFP is 57 percent higher in the large cities, and the high-skilled are 80 percent more productive relative to the middle-skilled types. The technology seems to more reasonably capture the TFP differences and productivity y_3 for high γ , corresponding to high elasticities of substitution. This is consistent with the fact that we do not condition on age (or other observables for that matter), where the age elasticity within skill group is typically large (of the order of 5).

In summary, quantitatively we find that the productivity difference between the top- and bottom-skilled workers is substantial and on the order

of at least four times bigger. For all specifications the degree of extreme-skill complementarity is positive, ranging from 0.8 percent to 4 percent depending on the returns to scale. In all specifications, the differences in TFP across large and small cities are big. Even for the lowest estimate with limited decreasing returns, TFP is about 60 percent higher in large cities.

IX. Conclusion

We propose a tractable theory of spatially dispersed production with perfectly mobile heterogeneous inputs, that is, skilled labor. Differences in TFP lead to differences in demand for skills across cities. In general equilibrium, wages and housing prices clear the labor and housing markets. Perfect mobility of citizens leads to utility equalization by skill. We show that cities with a higher TFP are larger and that a CES production technology entails identical skill distributions across cities with different productivity. We consider two alternative hypotheses concerning differential complementarities/substitutabilities between skills and derive the implication for the equilibrium skill distribution across cities. First, when there are complementarities between extreme skills, the firm size distribution in larger cities has thicker tails. Instead, when there are complementarities between the top skills, there is first-order stochastic dominance of the skill distribution in large cities.

We also find robust empirical evidence from US data for thick tails in the skill distribution. Adjusting wages for housing prices by means of a hedonic price index, we find that average skills are constant, but the standard deviation increases with city size. Big cities have big real inequality. Given the theory, this provides empirical support for the extreme-skill complementarity hypothesis: the productivity of the high-skilled is enhanced most by the providers of low-skilled services.

These findings contribute to our understanding of the urban wage premium. Not only do we establish robustly that higher wages are not due to higher average skill, but we also find that there is an urban inequality “premium.” In the presence of extreme-skill complementarities, this indicates that extreme skills multiply TFP differences.

Finally, our method and results can provide new insights into the role of complementarities in production. At the economywide level, we know remarkably little about the skill composition across firms of different sizes, for example, and even less about the pattern of complementarities between differentially skilled workers within firms. Understanding the patterns of complementarities not only is important for the efficient allocation of resources. As we have demonstrated in this paper, they are also key for the equitable distribution of the output of production.

Appendix A

Theory

Preliminaries

The full system of equations that pins down the equilibrium allocation can be written as, for $j \in \{1, 2\}$,

$$\begin{aligned}
 \lambda A_j (m_{3j}^\gamma y_3 + m_{1j}^\gamma y_1)^{\lambda-1} \gamma m_{1j}^{\gamma-1} y_1 - w_{1j} &= 0, \\
 \gamma A_j m_{2j}^{\gamma-1} y_2 - w_{2j} &= 0, \\
 \lambda A_j (m_{3j}^\gamma y_3 + m_{1j}^\gamma y_1)^{\lambda-1} \gamma m_{3j}^{\gamma-1} y_3 - w_{3j} &= 0, \\
 \frac{w_{i1}}{w_{i2}} &= \left(\frac{p_1}{p_2} \right)^\alpha \quad \forall i \in \{1, 2, 3\}, \\
 C_1 m_{i1} + C_2 m_{i2} &= M_i \quad \forall i \in \{1, 2, 3\}, \\
 h_{ij} &= \frac{\alpha w_{ij}}{p_j} \quad \forall i \in \{1, 2, 3\}, \\
 \sum_{i=1}^3 h_{ij} m_{ij} &= H.
 \end{aligned} \tag{A1}$$

Since $w_{i2} = (p_2/p_1)^\alpha w_{i1}$, we can equate the first-order conditions to obtain

$$\begin{aligned}
 A_1 (m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{11}^{\gamma-1} &= \left(\frac{p_1}{p_2} \right)^\alpha A_2 (m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{12}^{\gamma-1}, \\
 A_1 m_{21}^{\gamma-1} &= \left(\frac{p_1}{p_2} \right)^\alpha A_2 m_{22}^{\gamma-1}, \\
 A_1 (m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{31}^{\gamma-1} &= \left(\frac{p_1}{p_2} \right)^\alpha A_2 (m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{32}^{\gamma-1}, \\
 C_1 m_{i1} + C_2 m_{i2} &= M_i \quad \forall i \in \{1, 2, 3\}, \\
 m_{2j}^\gamma y_2 + \lambda (m_{3j}^\gamma y_3 + m_{1j}^\gamma y_1)^\lambda &= H \frac{p_j}{\alpha \gamma A_j}.
 \end{aligned} \tag{A2}$$

From the first and the third equations, we obtain

$$\frac{A_1 (m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{11}^{\gamma-1}}{A_1 (m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{31}^{\gamma-1}} = \frac{\left(\frac{p_1}{p_2} \right)^\alpha A_2 (m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{12}^{\gamma-1}}{\left(\frac{p_1}{p_2} \right)^\alpha A_2 (m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{32}^{\gamma-1}}. \tag{A3}$$

After rearranging and using market clearing ($C_1 m_{i1} + C_2 m_{i2} = M_i$), we can write this as

$$m_{i1} = \frac{M_i}{M_3} m_{31}, \tag{A4}$$

and we can write the first equation in (A2) as

$$\frac{(m_{31}^\gamma y_3 + m_{11}^\gamma y_1)^{\lambda-1} m_{21}^{\gamma-1}}{m_{21}^{\gamma-1}} = \frac{(m_{32}^\gamma y_3 + m_{12}^\gamma y_1)^{\lambda-1} m_{22}^{\gamma-1}}{m_{22}^{\gamma-1}}. \tag{A5}$$

Now given the symmetry assumption $m_{11} = m_{31}$ and $m_{12} = m_{32}$, this then implies

$$m_{21} = \left(\frac{m_{11}}{m_{12}}\right)^{(\lambda\gamma-1)/(\gamma-1)} m_{22}, \tag{A6}$$

and we substitute it back in (A2) and rearrange to get

$$\left(\frac{m_{11}}{m_{12}}\right)^{\lambda\gamma-1} = \left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}. \tag{A7}$$

Using the fact that

$$m_{12} = \frac{M_1}{C_2} - \frac{C_1}{C_2} m_{11},$$

we have

$$m_{11} = \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)} \frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}} \quad \text{and} \quad m_{12} = \frac{\frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}, \tag{A8}$$

and likewise for the other expressions for m_{2j} and m_{3j} :

$$m_{21} = \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\gamma-1)} \frac{M_2}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\gamma-1)}} \quad \text{and} \quad m_{22} = \frac{\frac{M_2}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\gamma-1)}}, \tag{A9}$$

$$m_{31} = \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)} \frac{M_3}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}} \quad \text{and} \quad m_{32} = \frac{\frac{M_3}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}. \tag{A10}$$

Substituting the m_{ij} 's in the last equation in (A2) and rearranging, we get

$$\begin{aligned} & \left[\left(\frac{p_2}{p_1} \right)^{[1-\lambda\gamma(1-\alpha)]/(1-\lambda\gamma)} \left(\frac{A_1}{A_2} \right)^{1/(1-\lambda\gamma)} - 1 \right] \lambda \left\{ \frac{\frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2} \right)^\alpha \frac{A_2}{A_1} \right]^{1/(\lambda\gamma-1)}}} \right\}^{\lambda\gamma} (y_3 + y_1)^\lambda \\ & + \left[\left(\frac{p_2}{p_1} \right)^{[1-\gamma(1-\alpha)]/(1-\gamma)} \left(\frac{A_1}{A_2} \right)^{1/(1-\gamma)} - 1 \right] \left\{ \frac{\frac{M_2}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2} \right)^\alpha \frac{A_2}{A_1} \right]^{1/(\gamma-1)}}} \right\}^\gamma y_2 \left\{ \frac{p_1}{p_2} \frac{A_2}{A_1} \right\} = 0, \end{aligned} \quad (*)$$

where we have used the fact that $\gamma < 1$, $\lambda > 1$, and $\lambda\gamma < 1$. We can now establish lemma 1.

LEMMA 1. Let $A_1 > A_2$, $\lambda\gamma < 1$, and $\gamma < 1$. Then housing prices in the more productive city are larger, $p_1 > p_2$.

Proof. In order to satisfy the equality (*), the only terms that can be negative are the ones in between brackets. Since $A_1/A_2 > 1$ and $\min\{1/(1-\lambda\gamma), 1/(1-\gamma)\} > 1$, the only way one of these terms is negative is if

$$\min \left\{ \left(\frac{p_2}{p_1} \right)^{[1-\lambda\gamma(1-\alpha)]/(1-\lambda\gamma)} \left(\frac{A_1}{A_2} \right)^{1/(1-\lambda\gamma)}, \left(\frac{p_2}{p_1} \right)^{[1-\gamma(1-\alpha)]/(1-\gamma)} \left(\frac{A_1}{A_2} \right)^{1/(1-\gamma)} \right\} < 1.$$

However, since $[1-\lambda\gamma(1-\alpha)]/(1-\lambda\gamma)$ and $[1-\gamma(1-\alpha)]/(1-\gamma)$ are positive, this is possible only if $p_2/p_1 < 1 \Rightarrow p_2 < p_1$. QED

Proof of Theorem 1

From lemma 1, we know that $p_1 > p_2$. Since $\lambda > 1$ and $\lambda\gamma < 1$, we have that

$$\frac{1}{1-\lambda\gamma} > \frac{1}{1-\gamma} > 1$$

and

$$\frac{1-\lambda\gamma(1-\alpha)}{1-\lambda\gamma} > \frac{1-\gamma(1-\alpha)}{1-\gamma} > 1.$$

Since we know that $A_1 > A_2$, we have that the first term in brackets in (*) is positive if

$$\frac{p_2}{p_1} > \left(\frac{A_2}{A_1} \right)^{1/[1-\lambda\gamma(1-\alpha)]} \quad (A11)$$

while the second term in brackets is positive if

$$\frac{p_2}{p_1} > \left(\frac{A_2}{A_1}\right)^{1/[1-\gamma(1-\alpha)]} \tag{A12}$$

Since $A_2/A_1 < 1$, we have that

$$\left(\frac{p_2}{p_1}\right)^\alpha \frac{A_1}{A_2} \in \left(\left(\frac{A_1}{A_2}\right)^{[(1-\lambda\gamma)(1-\alpha)]/[1-\lambda\gamma(1-\alpha)]}, \left(\frac{A_1}{A_2}\right)^{[(1-\gamma)(1-\alpha)]/[1-\gamma(1-\alpha)]} \right) > 1. \tag{A13}$$

From the expressions for m_{ij}

$$m_{11} = \frac{\left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)} \frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}, \tag{A14}$$

$$m_{12} = \frac{\frac{M_1}{C_2}}{1 + \frac{C_1}{C_2} \left[\left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1}\right]^{1/(\lambda\gamma-1)}}, \tag{A15}$$

and

$$\left(\frac{p_2}{p_1}\right)^\alpha \frac{A_1}{A_2} > 1,$$

we have that $m_{11} > m_{12}$, and likewise $m_{21} > m_{22}$ and $m_{31} > m_{32}$. Finally, since

$$S_j = m_{1j} + m_{2j} + m_{3j}, \tag{A16}$$

it immediately follows that $S_1 > S_2$. QED

Proof of Theorem 2

Consider the distributions, denoted by $pdf_{ij} = m_{ij}/S_j$, where we denote by

$$Z = \left(\frac{p_1}{p_2}\right)^\alpha \frac{A_2}{A_1} < 1.$$

Then we can write

$$pdf_{11} = \frac{\left[Z^{1/(\lambda\gamma-1)} \frac{M_1}{C_2} \right] / \left[1 + \frac{C_1}{C_2} Z^{1/(\lambda\gamma-1)} \right]}{\left\{ \left[Z^{1/(\lambda\gamma-1)} \frac{M_1 + M_3}{C_2} \right] / \left[1 + \frac{C_1}{C_2} Z^{1/(\lambda\gamma-1)} \right] \right\} + \left\{ \left[Z^{1/(\gamma-1)} \frac{M_2}{C_2} \right] / \left[1 + \frac{C_1}{C_2} Z^{1/(\gamma-1)} \right] \right\}}, \tag{A17}$$

$$pdf_{12} = \frac{\frac{M_1}{C_2} / \left[1 + \frac{C_1}{C_2} Z^{1/(\lambda\gamma-1)} \right]}{\left\{ \frac{M_1 + M_3}{C_2} / \left[1 + \frac{C_1}{C_2} Z^{1/(\lambda\gamma-1)} \right] \right\} + \left\{ \frac{M_2}{C_2} / \left[1 + \frac{C_1}{C_2} Z^{1/(\gamma-1)} \right] \right\}}. \tag{A18}$$

Then

$$\frac{pdf_{11}}{pdf_{12}} = \frac{Z^{1/(\lambda\gamma-1)} \{ (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + M_2 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \}}{Z^{1/(\lambda\gamma-1)} (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \times Z^{1/(\gamma-1)} M_2} > 1. \tag{A19}$$

Recall that $Z < 1$, and therefore,

$$Z^{1/(\lambda\gamma-1)} = \left(\frac{1}{Z} \right)^{1/(1-\lambda\gamma)};$$

then since $1/Z > 1$, the larger the exponent, the larger is $(1/Z)^{1/(1-\lambda\gamma)}$. Since $\lambda > 1$, $\lambda\gamma < 1$, and $1/(1 - \lambda\gamma) > 1/(1 - \lambda)$, it follows that $pdf_{11} > pdf_{12}$.

Similarly, we can show that $pdf_{31} > pdf_{32}$:

$$\frac{pdf_{31}}{pdf_{32}} = \frac{Z^{1/(\lambda\gamma-1)} \{ (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + M_2 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \}}{Z^{1/(\lambda\gamma-1)} (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \times Z^{1/(\gamma-1)} M_2} > 1. \tag{A20}$$

Finally,

$$pdf_{21} = \frac{[C_2 + C_1 Z^{1/(\lambda\gamma-1)}] Z^{1/(\gamma-1)} M_2}{Z^{1/(\lambda\gamma-1)} (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\lambda-1)}] + Z^{1/(\gamma-1)} M_2 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]},$$

$$pdf_{22} = \frac{[C_2 + C_1 Z^{1/(\lambda\gamma-1)}] M_2}{(M_1 + M_3) [C_2 + C_1 Z^{1/(\gamma-1)}] + M_2 [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}.$$

Then

$$\frac{pdf_{21}}{pdf_{22}} = \frac{Z^{1/(\gamma-1)} (M_1 + M_3) [C_2 + C_1 Z^{1/(\lambda-1)}] + M_2 [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}{Z^{1/(\lambda\gamma-1)} (M_1 + M_3) \times [C_2 + C_1 Z^{1/(\lambda-1)}] + Z^{1/(\lambda-1)} M_2 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]} \tag{A21}$$

Again, since $Z^{1/(\gamma-1)} < Z^{1/(\lambda\gamma-1)}$, now we have that $pdf_{21} < pdf_{22}$. QED

Proof of Theorem 3

The proof follows closely the logic of theorem 2:

$$pdf_{11} = \frac{[C_2 + C_1 Z^{1/(\lambda\gamma-1)}] Z^{1/(\gamma-1)} M_1}{Z^{1/(\lambda\gamma-1)} (M_2 + M_3) \times [C_2 + C_1 Z^{1/(\lambda-1)}] + Z^{1/(\gamma-1)} M_1 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]},$$

$$pdf_{12} = \frac{[C_2 + C_1 Z^{1/(\lambda\gamma-1)}] M_1}{(M_2 + M_3) [C_2 + C_1 Z^{1/(\gamma-1)}] + M_1 [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}.$$

Then

$$\frac{pdf_{11}}{pdf_{12}} = \frac{Z^{1/(\gamma-1)} (M_2 + M_3) [C_2 + C_1 Z^{1/(\lambda-1)}] + M_1 [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}{Z^{1/(\lambda\gamma-1)} (M_2 + M_3) \times [C_2 + C_1 Z^{1/(\lambda-1)}] + Z^{1/(\lambda-1)} M_1 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]} \tag{A22}$$

$$< 1,$$

since $Z^{1/(\gamma-1)} < Z^{1/(\lambda\gamma-1)}$ and

$$pdf_{21} = \frac{[C_2 + C_1 Z^{1/(\gamma-1)}] Z^{1/(\lambda\gamma-1)} M_2}{Z^{1/(\lambda\gamma-1)} (M_2 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \times Z^{1/(\gamma-1)} M_1},$$

$$pdf_{22} = \frac{[C_2 + C_1 Z^{1/(\gamma-1)}] M_2}{(M_2 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + M_1 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}.$$

Then

$$\frac{pdf_{21}}{pdf_{22}} = \frac{Z^{1/(\lambda\gamma-1)}(M_2 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + M_1 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}{Z^{1/(\lambda\gamma-1)}(M_2 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \times Z^{1/(\gamma-1)}M_1} \quad (\text{A23})$$

$$> 1.$$

Finally, analogously to pdf_{21}/pdf_{22} , we can derive pdf_{31}/pdf_{32} :

$$\frac{pdf_{31}}{pdf_{32}} = \frac{Z^{1/(\lambda\gamma-1)}(M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + M_2 \times [C_2 + C_1 Z^{1/(\lambda\gamma-1)}]}{Z^{1/(\lambda\gamma-1)}(M_1 + M_3) \times [C_2 + C_1 Z^{1/(\gamma-1)}] + [C_2 + C_1 Z^{1/(\lambda\gamma-1)}] \times Z^{1/(\gamma-1)}M_2} \quad (\text{A24})$$

$$> 1.$$

QED

Quantifying the Production Technology: Derivation of the System

The equilibrium allocation can explicitly be represented by the following system of five equations in five unknowns: λ , A_1 , A_2 , y_1 , and y_3 , where y_2 is normalized to one:

$$\lambda = \frac{1}{\gamma} \left[1 + (\gamma - 1) \frac{\log\left(\frac{C_2 m_{21}}{M_2 - C_1 m_{21}}\right)}{\log\left(\frac{C_2 m_{11}}{M_1 - C_1 m_{11}}\right)} \right],$$

$$A_1 = \frac{w_{21}}{\gamma y_2 m_{21}^{\gamma-1}},$$

$$A_2 = A_1 \left(\frac{p_2}{p_1}\right)^\alpha \left(\frac{C_2 m_{21}}{M_2 - C_1 m_{21}}\right)^{\gamma-1}, \quad (\text{A25})$$

$$y_1 = \left\{ \frac{w_{11}}{\lambda \gamma A_1 [m_{11} + m_{31} (w_{31}/w_{11})]^{\lambda-1} m_{11}^{\lambda(\gamma-1)}} \right\}^{1/\lambda},$$

$$y_3 = \left\{ \frac{w_{31}}{\lambda \gamma A_1 [m_{31} + m_{11} (w_{11}/w_{31})]^{\lambda-1} m_{31}^{\lambda(\gamma-1)}} \right\}^{1/\lambda}.$$

From equation (A9), we can solve for Z :

$$Z = \left(\frac{C_2 m_{21}}{M_2 - C_1 m_{21}} \right)^{\gamma-1}. \quad (\text{A26})$$

Now solve the expression for m_{11} in terms of λ , and substitute for this expression for Z :

$$\lambda = \frac{1}{\gamma} \left[1 + (\gamma - 1) \frac{\log\left(\frac{C_2 m_{21}}{M_2 - C_1 m_{21}}\right)}{\log\left(\frac{C_2 m_{11}}{M_1 - C_1 m_{11}}\right)} \right]. \tag{A27}$$

Therefore, the solution for λ , Z is simply the explicit solution to the equations (A26) and (A27).

Now from the second first-order condition (eq. [10]), we obtain

$$y_2 = \frac{w_{21}}{\gamma A_1 m_{21}^{\gamma-1}}, \tag{A28}$$

while from the first (eq. [9]) and the third first-order conditions (eq. [11]), we can solve explicitly for y_1 and y_3 :

$$\begin{aligned} y_1 &= \frac{w_{11}}{w_{31}} \frac{m_{31}^{\gamma-1}}{m_{11}^{\gamma-1}} \left\{ \frac{w_{31}}{\lambda \gamma A_1 [m_{31} + m_{11}(w_{11}/w_{31})]^{\lambda-1} m_{31}^{\lambda(\gamma-1)}} \right\}^{1/\lambda} \\ &= \left\{ \frac{w_{11}}{\lambda \gamma A_1 [m_{11} + m_{31}(w_{31}/w_{11})]^{\lambda-1} m_{11}^{\lambda(\gamma-1)}} \right\}^{1/\lambda}, \end{aligned} \tag{A29}$$

$$y_3 = \left\{ \frac{w_{31}}{\lambda \gamma A_1 [m_{31} + m_{11}(w_{11}/w_{31})]^{\lambda-1} m_{31}^{\lambda(\gamma-1)}} \right\}^{1/\lambda}. \tag{A30}$$

Finally, we obtain the expression for A_1 from the second first-order condition (10) and the expression for A_2 from substituting for $Z = (p_1/p_2)^\alpha (A_2/A_1)$ in (A26):

$$A_1 = \frac{w_{21}}{\gamma y_2 m_{21}^{\gamma-1}}, \tag{A31}$$

$$A_2 = A_1 \left(\frac{p_2}{p_1} \right)^\alpha \left(\frac{C_2 m_{21}}{M_2 - C_1 m_{21}} \right)^{\gamma-1}. \tag{A32}$$

Normalizing $y_2 = 1$, equations (A27), (A31), (A32), (A29), and (A30) constitute the system (A25).

Appendix B

Data

Wage Data

Wage data are taken from the CPS, a joint effort between the US Bureau of Labor Statistics (BLS) and the Census Bureau (<http://www.bls.gov/cps/>). The CPS is a monthly survey and is used by the US government to calculate the official unemployment and labor force participation figures. We used the 2009 merged outgoing rotation groups (MORG) as provided by the NBER.²⁵ The MORG are extracts of the basic monthly data during the household's fourth and eighth months in the survey, when usual weekly hours/earnings are asked.

We use the variable *earnwke* as created by the NBER.²⁶ This variable reports earnings per week in the current job. It includes overtime, tips, and commissions. For hourly workers, item 25a ("How many hours per week does . . . usually work at this job?") times item 25c ("How much does . . . earn per hour?") appears here. For weekly workers, item 25d ("How much does . . . usually earn per week at this job before deductions?") appears here.

The NBER version of the CPS identifies the CBSA of the observation. It reports the New England city and town areas definition and codes for metro areas in the six New England states and the Federal Information Processing Standards definition and codes for all other states. Table B1 shows the 10 largest and 10 smallest included CBSAs.

We restrict the sample to full-time workers (between 36 and 60 usual hours per week). We also drop the lowest 0.5 percent of wages as a pragmatic way of eliminating likely misreported wages close to zero. Our final wage sample includes 76,821 workers in 254 identified CBSAs out of the 320,941 surveyed persons. CPS wage data are in 2009 top-coded at a weekly wage of \$2,884.61, which applies to 2,308 or 3.0 percent of workers. All estimations use the weights in the variable *earnwt* provided by the NBER.

Our baseline results use CPS wage data because they are generally considered of higher quality than census data (see, e.g., Baum-Snow and Neal 2009). However, the CPS has two disadvantages: it has relatively low top codes and it does not identify the location of the household within cities. We therefore alternatively use wage data from the 2009 ACS collected by the US Census Bureau.²⁷ The data are provided by the Minnesota Population Center in its Integrated Public Use Microdata Series (IPUMS).²⁸

We use the variable *incwage*, which measures yearly wage and salary income. We restrict our sample to full-time (between 36 and 60 usual hours per week) and full-year (between 48 and 52 weeks per year) workers. The yearly wage is divided by the number of weeks worked to get weekly wages comparable to the CPS data. ACS

²⁵ The Stata data file is available at <http://www.nber.org/morg/annual/morg09.dta>.

²⁶ See details of the variable creation at the NBER website, <http://www.nber.org/cps/>.

²⁷ See <http://www.census.gov/acs/www/> for more information on the survey.

²⁸ See Ruggles et al. (2010) for the data source and <http://usa.ipums.org/usa/> for a detailed description of data and variables.

TABLE B1
RANK OF CITIES BY 2009 POPULATION

Rank	City	Population
1	New York–Northern New Jersey–Long Island, NY-NJ-PA	19,069,796
2	Los Angeles–Long Beach–Santa Ana, CA	12,874,797
3	Chicago–Naperville–Joliet, IL-IN-WI	9,580,567
4	Dallas–Fort Worth–Arlington, TX	6,447,615
5	Philadelphia–Camden–Wilmington, PA-NJ-DE-MD	5,968,252
6	Houston–Sugar Land–Baytown, TX	5,867,489
7	Miami–Fort Lauderdale–Pompano Beach, FL	5,547,051
8	Washington–Arlington–Alexandria, DC-VA-MD-WV	5,476,241
9	Atlanta–Sandy Springs–Marietta, GA	5,475,213
10	Boston–Cambridge–Quincy, MA-NH	4,588,680
...		
245	Farmington, NM	124,131
246	Bowling Green, KY	120,595
247	Harrisonburg, VA	120,271
248	Lawrence, KS	116,383
249	Victoria, TX	115,396
250	Anniston–Oxford, AL	114,081
251	Lawton, OK	113,228
252	Kankakee–Bradley, IL	113,215
253	Michigan City–La Porte, IN	111,063
254	Decatur, IL	108,204

NOTE.—Cities are defined as core-based statistical areas. The Office of Management and Budget defines 940 metropolitan and micropolitan areas, of which we use the ones with population above 100,000 and where housing prices are observed.

data are top-coded at the 99.5th percentile of each state. We also drop the lowest 0.1 percent of wages as a pragmatic way of eliminating likely misreported wages close to zero. Our final ACS wage sample includes 654,043 workers in 293 identified CBSAs (with population above 100,000).

The ACS discloses the so-called PUMAs, which are areas with a maximum of 179,405 housing units and only partly overlap with political borders of towns and counties. We use the Geographic Correspondence Engine with Census 2000 Geography from the Missouri Census Data Center (MCDC; <http://mcdc.missouri.edu/websas/geocorr2k.html>) to link PUMAs to CBSAs. The MCDC data match every urban PUMA code to one or more CBSA codes and report the fraction of housing units that are matched. We assign a PUMA to a CBSA if this fraction is bigger than 33 percent. In cases in which the PUMA does not fully belong to a CBSA, we assign the PUMA to the CBSA to which most of its housing units belong. Our final sample contains data from 533 metropolitan or micropolitan CBSAs out of a total of 940 existing CBSAs. Note that we do not use the metropolitan area code provided in the ACS in the variable *metaread*. This variable reports a mixture of metropolitan area codes (MSA, primary MSA, central city or county), which is difficult to match with the CBSA definition. Our final sample contains 273,761 rental units in 533 CBSAs (with population above 100,000) and 1,884 PUMAs.

Local Housing and Commodity Price Indices

We use the 2009 ACS for our baseline housing price estimates.

The variable *rent* reports the monthly contract rent for rental units in contemporary dollars. We also use all the reported housing characteristics of the unit: *rooms* is the number of rooms, *unitsstr* is the units in the structure (in eight groups), and *builtyr* is the age of the structure (in 13 age groups).

We drop housing units in group quarters and farmhouses; drop mobile homes, trailers, boats, and tents; and use data from housing units only in identified metropolitan or micropolitan CBSAs.

For robustness checks, we also purchased the ACCRA Cost of Living Index from C2ER. ACCRA data are collected by local chambers of commerce and similar organizations that have volunteered to participate. They are reported for 302 CBSAs and 23 metropolitan divisions for the 11 largest CBSAs. The ACCRA Cost of Living Index consists of six major categories: grocery items, housing, utilities, transportation, health care, and miscellaneous goods and services. These major categories in turn are composed of subcategories, each of which is represented by one or more items in the index. In total, local prices of 60 items are reported, for example, T-bone steak (item 1), phone (31), gasoline (33), Lipitor (38), pizza (40) haircut (42), and movie (52). Indices for major categories and an overall composite index are calculated as weighted averages, where weights come from the CEX conducted by the BLS. We use the average of quarterly data from Q2.2008 to Q2.2009 in order to minimize the number of missing cities from nonreporting places. We use the average across metropolitan divisions to match ACCRA data to our wage data.

Hedonic Regression to Calculate Housing Price Index

We model housing as a homogeneous good h with a location-specific per unit price p . In practice, however, housing differs in many observable dimensions. Observed housing prices therefore reflect both the location and the physical characteristics of the unit. Sieg et al. (2002) show the conditions under which housing can be treated as if it were homogeneous and how to construct a price index for it. Take our Cobb-Douglas utility function,

$$u(c, h(z)) = c^{1-\alpha} h^\alpha(z), \quad (\text{B1})$$

and assume that housing $h(z)$ is a function, for simplicity of exposition only, of two characteristics $z = (z_1, z_2)$ with a nested Cobb-Douglas structure

$$h(z) = z_1^\delta z_2^{1-\delta}. \quad (\text{B2})$$

The indirect utility given the market prices q_1 and q_2 for, respectively, characteristics z_1 and z_2 is then

$$U_i = \alpha^\alpha (1 - \alpha)^{1-\alpha} (L q_1^\delta q_2^{1-\delta})^{-\alpha} w, \quad (\text{B3})$$

where $L = 1/[\delta^\delta (1 - \delta)^{1-\delta}]$. Defining the price index $p = L q_1^\delta q_2^{1-\delta}$, the indirect utility is

$$U_i = \alpha^\alpha (1 - \alpha)^{1-\alpha} \frac{w}{p^\alpha} \tag{B4}$$

and thus is identical to the one derived assuming homogeneous housing h with market price p . The subexpenditure function $e(q_1, q_2, h)$ is defined as the minimum expenditure necessary to obtain h units of housing and is given by

$$e(q_1, q_2, h) = Lq_1^\delta q_2^{1-\delta} h = ph = pz_1^\delta z_2^{1-\delta}. \tag{B5}$$

Taking logarithms and assuming that we observe z_1 but not z_2 yields a linear hedonic regression model

$$\log(e_n) = \log(p_i) + \delta \log(z_{1jn}) + u_n, \tag{B6}$$

where e_n is the observed rental price of housing unit n and $\log(p_i)$ is a city-specific common component of housing prices. We can therefore estimate the city-specific price level as a location-specific fixed effect in a simple hedonic regression of log rental prices on the physical characteristics.

Table B2 shows the results of the hedonic regressions for rental units using data from the 2009 ACS. Column 1 shows the results with 533 fixed effects for cities (CBSAs) and column 2 with 1,844 fixed effects for neighborhoods (PUMAs). We use all relevant housing characteristics in the data and add all categories as dummy variables without functional form assumptions. All coefficients are highly significant with expected signs: housing prices increase with the number of rooms and decrease with the age of the structure. We find a nonmonotonic relationship in the numbers of units in the structure with highest prices for single-family detached homes and lowest prices for three to four-family buildings.

We standardize the housing price index such that the weighted (by housing units) average equals one. Table B3 shows the resulting housing price indices for the highest- and lowest-priced cities in our sample. The highest-priced city is San Jose, California, with rental prices 74 percent above average urban prices; the lowest-priced city with more than 100,000 inhabitants is Anniston, Alabama, with prices 64 percent below average.

Expenditure Data for Estimation of the Nonlinear Engel Curve

We use public use microdata (PUMD) from the CEX, which are provided by the BLS.²⁹ Our sample consists of interview data for 4,581 households in the first quarter of 2011 in the data file fmli111x. The PUMD reports the exact city (MSA) for 2,021 households living in one of the 21 largest MSAs. We merge these data with our hedonic price indices from the corresponding CBSAs. The housing expenditure shares, s_n , are measured as “expenditures on shelter this quarter” (vari-

²⁹ See <http://www.bls.gov/ceX/> for details on the CEX and its public use individual data version. The CEX allows us to calculate expenditure share as a ratio of total expenditures. The ACS does not provide total expenditures of individuals or households. Housing shares as a ratio of reported income are extremely noisy, with housing expenditure shares above 100 percent for a large fraction of low-income households. We did not get any reasonable Stone-Geary parameters based on ACS data.

TABLE B2
HEDONIC REGRESSIONS FOR RENTAL UNITS

	CBSA Level	PUMA Level
Number of rooms:		
1	-.2314*** (.0056)	-.2238*** (.0055)
2	-.1658*** (.0050)	-.1863*** (.0049)
3	-.1329*** (.0031)	-.1386*** (.0030)
4	0	0
5	.0760*** (.0033)	.0798*** (.0031)
6	.1614*** (.0041)	.1592*** (.0039)
7	.2405*** (.0057)	.2313*** (.0055)
8	.2877*** (.0077)	.2717*** (.0074)
9+	.3341*** (.0082)	.3049*** (.0079)
Age of structure:		
1939 or earlier	-.3068*** (.0053)	-.2700*** (.0053)
1940–49	-.3603*** (.0062)	-.3219*** (.0061)
1950–59	-.3167*** (.0055)	-.2970*** (.0054)
1960–69	-.2887*** (.0053)	-.2793*** (.0052)
1970–79	-.2553*** (.0050)	-.2542*** (.0049)
1980–89	-.1758*** (.0052)	-.1838*** (.0050)
1990–99	-.0780*** (.0054)	-.0838*** (.0052)
2000–2004	0	0
2005	.0122 (.0097)	.0223** (.0094)
2006	.0421*** (.0099)	.0537*** (.0095)
2007	.0548*** (.0104)	.0621*** (.0100)
2008	.1029*** (.0135)	.1139*** (.0130)
2009	.0343 (.0444)	.0347 (.0427)
Units in structure:		
1-family house detached	0	0
1-family house attached	-.0635*** (.0050)	-.0677*** (.0049)
2-family building	-.1257*** (.0045)	-.1289*** (.0044)
3–4 family building	-.1314*** (.0042)	-.1434*** (.0041)
5–9 family building	-.1239*** (.0042)	-.1532*** (.0041)
10–19 family building	-.0786*** (.0043)	-.1171*** (.0042)
20–49 family building	-.1023*** (.0048)	-.1354*** (.0047)
50+ family building	-.0929*** (.0045)	-.1413*** (.0045)
Constant	6.5728*** (.0481)	6.0277*** (.0523)
CBSA fixed effects	Yes	No
PUMA fixed effects	No	Yes
Observations (rental units)	273,761	273,761
Number of CBSAs	533	533
Number of PUMA regions		1,884

NOTE.—Reference groups are indicated by 0. Standard errors are in parentheses.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

able SHELTCQ) divided by total expenditures this quarter (TOTEXPCQ). Weekly wages, w_t , are measured as “income before taxes in past 12 months” (FINCBTAX) divided by 52. Housing prices, p_t , are measured as the hedonic price index using the 2009 ACS. We restrict our sample to weekly wage incomes above \$175 because very low wages become extremely large inverse values, which would almost fully determine the regression. The limit of \$175 is also the lower bound in our CPS wage data (see the data section on wages) and guarantees that the indirect

TABLE B3
RANK OF CITIES BY ESTIMATED HOUSING PRICE INDEX

Rank	City	Population	Rent Index
1	San Jose–Sunnyvale–Santa Clara, CA	1,839,700	1.74
2	San Francisco–Oakland–Fremont, CA	4,317,853	1.64
3	Santa Barbara–Santa Maria–Goleta, CA	407,057	1.62
4	Oxnard–Thousand Oaks–Ventura, CA	802,983	1.62
5	Honolulu, HI	907,574	1.61
6	Los Angeles–Long Beach–Santa Ana, CA	12,874,797	1.55
7	San Diego–Carlsbad–San Marcos, CA	3,053,793	1.51
8	Washington–Arlington–Alexandria, DC-VA-MD-WV	5,476,241	1.46
9	Napa, CA	134,650	1.43
10	Santa Cruz–Watsonville, CA	256,218	1.43
11	New York–Northern New Jersey–Long Island, NY-NJ-PA	19,069,796	1.41
...			
245	McAllen–Edinburg–Mission, TX	741,152	.50
246	Lawton, OK	113,228	.50
247	Lake Charles, LA	194,138	.49
248	Huntington–Ashland, WV-KY-OH	285,624	.49
249	Monroe, LA	174,086	.47
250	Johnstown, PA	143,998	.47
251	Brownsville–Harlingen, TX	396,371	.47
252	Decatur, AL	151,399	.46
253	Joplin, MO	174,300	.45
254	Anniston–Oxford, AL	114,081	.36

NOTE.—Housing price indices are based on hedonic regressions using the 2009 ACS.

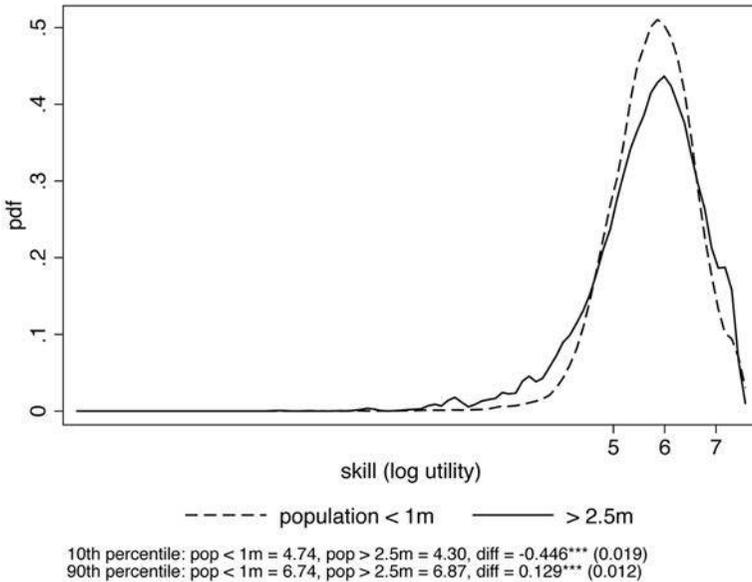


FIG. B1.—Expenditure share using Stone-Geary preferences, $\alpha = 0.22$, $\underline{h} = 250$

utility is defined for all observations. The regression with 1,569 observations results in $\hat{\alpha} = 0.224$ (SE = 0.005), $\hat{\beta} = 21.5$ (SE = 3.05), and $\hat{h} = 27.7$ (SE 3.79 using the delta method). As a robustness check, we also use values for \hat{h} different from the estimated value. As an extreme example, for $\hat{h} = 250$, we find that even though the shape of the distribution changes dramatically—especially for the low-income workers—the thick tails are still prominently present, as illustrated in figure B1.

Appendix C

Preferences for Services and Home Production

In this appendix we explore the microfoundations for the complementarities between extreme skills. We consider an explanation based on preferences for low-skilled services in combination with home production. Households have preferences for low-skilled services. Those services can be produced with home production, and they can be traded on the market.

Citizens have preferences over the quantity of a consumption good c , services s , and the amount of housing h represented by $u(h, s, c) = h^\alpha s^\beta c^{1-\alpha-\beta}$, where α, β , and $(\alpha + \beta) \in [0, 1]$. All workers have a unit endowment of time to be divided between home production denoted by t and market production $1 - t$. Home productivity is independent of the worker's skills. The amount of services generated in home production is equal to Γt^δ , where Γ is a positive parameter, but the agent incurs a quasi fixed cost K , a cost that is incurred only if there is positive production. Services can be traded at a city-specific price $r_j \geq 0$. As before, output in the formal sector generates a wage that is contingent on the worker's skill and produced with the technology $A_j F(\cdot)$ that either is CES or satisfies top- or extreme-skill complementarity.

A citizen i in city j chooses the bundle $\{h, s, c\}$ to maximize utility subject to the budget constraint:

$$\begin{aligned} \max_{\{h, s, c, t\}} u(c, h, s) &= h^\alpha s^\beta c^{1-\alpha-\beta} \\ \text{subject to } ph + rs + c &\leq w(1 - t) + I, \end{aligned} \quad (\text{C1})$$

where the quantities h, s, c, t , and w are all city and skill specific and the prices p and r are city specific; we have dropped the subscripts for notational convenience. We denote by I the income generated through services provided. The decision problem to produce services t is given by $\max_{0 < t \leq 1} \{r\Gamma t^\delta - K + wt\}$ if $t > 0$ and zero otherwise. Observe that K acts as a cost of entry in the services sector. Solving this problem, and given the cost K , the optimal solution satisfies

$$I = \max_i \left\{ r\Gamma \left(\frac{r\delta\Gamma}{w} \right)^{\delta/(1-\delta)} - K, 0 \right\}. \quad (\text{C2})$$

Because of the fixed cost, there will be an occupational choice decision whether or not to produce services, and if so, how much (t). With this optimal allocation of time between production of services and of market output, we can pin down the demands for consumption, services, and housing as

$$\begin{aligned}
 h &= \frac{\alpha}{p} [w(1 - t^*) + I], \\
 s &= \frac{\beta}{r} [w(1 - t^*) + I], \\
 c &= (1 - \alpha - \beta) [w(1 - t^*) + I],
 \end{aligned}
 \tag{C3}$$

where $t^* = \max\{(r\delta\Gamma/w)^{1/(1-\delta)}, 0\}$. Observe that the time spent in household production depends on the ratio r/w . The higher the wage, the less time the worker spends producing household services, and the more likely she is a net demander of services. Moreover, because of the cost K , those with high enough wages will choose not to produce services at all. Finally, the market-clearing conditions (including in the market for services) pin down equilibrium prices and close the model.

We cannot solve the model analytically. We have therefore performed various quantitative exercises to get an idea of the properties of the model. In particular, we want to find parameter configurations under which we obtain thick tails. We use the parameters $\alpha = 0.24$, $\beta = 0.2$, $\delta = 0.3$, $\gamma = 0.8$, $\Gamma = 1$, and $K = 0.2$, and from the data, $M_1 = 21,644,289$, $M_2 = 62,544,627$, $M_3 = 21,644,289$, $C_1 = 21$, $C_2 = 204$, and $H = 110,016.5$, which is based on the observed mean number of housing units. Then we use the technology parameters generated by the exercise in Section VIII. The objective is to obtain distributional properties that are consistent with thick tails.

First, we can robustly reproduce the thick-tails results in this model with services whenever there are extreme-skill complementarities in the market technology. This confirms that our approach in the baseline model is robust to the introduction of low-skilled services. At the same time, this may not be all that surprising especially whenever the mechanism of extreme-skill complementarities is strong enough. More challenging is whether thick tails obtain without extreme-skill complementarities. Our second finding is that with a CES market technology and without the quasi fixed costs K , the distributions with services are identical across cities. With $K > 0$ and CES, the distributions differ across cities, which results in first-order stochastic dominance of the small city: there are relatively more low-skilled workers in the large city, relatively more middle-skilled workers in the small city, and the same density of high-skilled workers in both cities; the average skills are lower in the large cities. Third, the model generates thick tails if preferences for services are combined with top-skill complementarity. The logic is that the top-skill complementarity generates the fat upper tail, whereas the demand for services generates the fat lower tail.

Table C1 shows this third result in a quantitative exercise. To obtain thick tails, the income share of services β must be large enough. When it drops below 20 percent, the lower tail disappears.³⁰ This means that households must spend nearly as much on low-skilled services as on housing. Second, the fixed cost K must also be large enough ($K = 0.2$ is a big barrier), which effectively inflates the price of services. Observe that there is no such entry barrier in the formal sector, where one would expect those to be at least as big. Finally, the top-skill complementarity must not be too strong.

³⁰ In the CEX, the direct expenditure on low-skilled services is around 5 percent.

TABLE C1
PREFERENCES FOR SERVICES AND HOME PRODUCTION

γ	ASSUMED PRODUCTION TECHNOLOGY							
	λ	A_1	A_2	y_1	y_2	y_3	C_1	C_2
.8	1.0193	19,118	9,065	.3189	1	1.4733	21	204
CITY j	OBSERVED MODEL OUTCOMES							
	w_{1j}	w_{2j}	w_{3j}	pdf_{1j}	pdf_{2j}	pdf_{3j}	p_j	r_j
1	430	1,315	2,330	.2128	.5693	.2178	4,473.43	653.88
2	261	774	1,419	.2016	.5986	.1998	763.33	372.82

Our simulations suggest that demand for services can indeed contribute to the fat lower tail and that it is broadly consistent with extreme-skill complementarity. To generate thick tails without extreme-skill complementarity, however, the required parameters are somewhat extreme. In particular, it requires a strong nonhomotheticity in the services technology, and the expenditure share on services must be unrealistically high. Many low-skilled services are demanded indirectly because they are inputs in production: cooks in restaurants, administrative staff in firms, and so forth. We find it therefore justified to model their role through the production technology. Finally, a prime advantage of our basic model with complementarities in the production technology is that we can solve and study it analytically.

References

- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4B, edited by David Card and Orley Ashenfelter, 1043–1171. Amsterdam: Elsevier.
- Albouy, David. 2008. "Are Big Cities Really Bad Places to Live? Estimating Quality of Life across Metropolitan Areas." Working Paper no. 14472, NBER, Cambridge, MA.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market." *A.E.R.* 103 (5): 1553–97.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Investigation." *Q.J.E.* 118 (4): 1279–1333.
- Bacolod, Marigee, Bernardo S. Blum, and William C. Strange. 2009. "Skills in the City." *J. Urban Econ.* 65 (2): 136–53.
- Baum-Snow, Nathaniel, and Derek Neal. 2009. "Mismeasurement of Usual Hours Worked in the Census and ACS." *Econ. Letters* 102 (1): 39–41.
- Baum-Snow, Nathaniel, and Ronni Pavan. 2012. "Understanding the City Size Wage Gap." *Rev. Econ. Studies* 79 (1): 88–127.
- Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2014. "Productive Cities: Sorting, Selection, and Agglomeration." *J.P.E.* 122 (3): 507–53.
- Blinder, Alan. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *J. Human Resources* 8:436–55.

- Broda, Christian, Ephraim Leibtag, and David E. Weinstein. 2009. "The Role of Prices in Measuring the Poor's Living Standards." *J. Econ. Perspectives* 23 (2): 77–97.
- Broda, Christian, and David E. Weinstein. 2008. "Understanding International Price Differences Using Barcode Data." Working Paper no. 14017 (May), NBER, Cambridge, MA.
- Chernozhukov, Victor, Ivan Fernández-Val, and Blaise Melly. 2013. "Inference on Counterfactual Distributions." *Econometrica* 81 (6): 2205–68.
- Davis, Morris A., and François Ortalo-Magné. 2011. "Household Expenditures, Wages, Rents." *Rev. Econ. Dynamics* 14 (2): 248–61.
- Davis, Steven J., R. Jason Faberman, and John Haltiwanger. 2006. "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links." *J. Econ. Perspectives* 20 (3): 3–26.
- De la Roca, Jorge, and Diego Puga. 2012. "Learning by Working in Big Cities." Manuscript, Centro de Estudios Monetarios y Financieros, Madrid.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64:1001–44.
- Eeckhout, Jan. 2004. "Gibrat's Law for (All) Cities." *A.E.R.* 94 (5): 1429–51.
- Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77 (3): 953–73.
- Fortin, Nicole M., Thomas Lemieux, and Sergio Firpo. 2011. "Decomposition Methods in Economics." In *Handbook of Labor Economics*, vol. 4A, edited by Orley Ashenfelter and David Card, 1–102. Amsterdam: Elsevier.
- Fu, S., and Stephen Ross. 2010. "Wage Premia in Employment Clusters: Agglomeration or Worker Heterogeneity?" Working Paper no. 10-04, Center Econ. Studies, US Census Bureau, Washington, DC.
- Fujita, Masahisa. 1989. *Urban Economic Theory: Land Use and City Size*. Cambridge: Cambridge Univ. Press.
- Gabaix, Xavier. 1999. "Zipf's Law for Cities: An Explanation." *Q.J.E.* 114 (3): 738–67.
- Gautier, Pieter A., M. Svarer, and C. N. Teulings. 2009. "Marriage and the City: Search Frictions and Sorting of Singles." *J. Urban Econ.* 67 (2): 206–18.
- Gautier, Pieter A., and C. N. Teulings. 2009. "Search and the City." *Regional Sci. and Urban Econ.* 39:251–65.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2009. "Job Polarization in Europe." *A.E.R. Papers and Proc.* 99 (2): 58–63.
- Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2011. "Endogenous Gentrification and Housing Price Dynamics." Manuscript, Univ. Chicago.
- Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan. 2003. "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation." *J.P.E.* 111 (3): 465–97.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *J.P.E.* 101 (3): 410–42.
- Kelso, Alexander S. J., and V. P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50:1483–1504.
- Koo, Jahyeong, Keith R. Phillips, and Fiona D. Sigalla. 2000. "Measuring Regional Cost of Living." *J. Bus. and Econ. Statis.* 18 (1): 127–36.
- Krusell, Per, Lee E. Ohanian, José-Victor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–54.

- Lucas, Robert E., Jr., and Esteban Rossi-Hansberg. 2002. "The Internal Structure of Cities." *Econometrica* 70 (4): 1445–76.
- Marshall, Alfred. 1890. *Principles of Economics*. London: Macmillan.
- Moretti, Enrico. 2013. "Real Wage Inequality." *American Econ. J.: Appl. Econ.* 5 (1): 65–103.
- Munshi, Kaivan. 2003. "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market." *Q.J.E.* 118 (2): 549–97.
- Oaxaca, Ronald. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *Internat. Econ. Rev.* 14:693–709.
- Rothe, Christoph. 2012. "Decomposing the Composition Effect: Understanding the Role of Covariates in Determining Between-Group Differences in Economic Outcomes." Manuscript, Columbia Univ.
- Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. 2010. "Integrated Public Use Microdata Series: Version 5.0" [machine-readable database]. Minneapolis: Univ. Minnesota.
- Schmidheiny, Kurt. 2006. "Income Segregation from Local Income Taxation When Households Differ in Both Preferences and Incomes." *Regional Sci. and Urban Econ.* 36 (2): 270–99.
- Sieg, Holger, V. Kerry Smith, H. Spencer Banzhaf, and Randy Walsh. 2002. "Interjurisdictional Housing Prices in Locational Equilibrium." *J. Urban Econ.* 52: 131–53.
- Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill. 2010. "Why Has House Price Dispersion Gone Up?" *Rev. Econ. Studies* 77 (4): 1567–1606.