Local labor markets

Jan Eeckhout

CREi

Centre de Recerca
en Economia Internacional

Local Labor Markets

Jan Eeckhout

1. Introduction

Labor markets are geographically differentiated. And while differentiated goods markets are connected through the substitution of one differentiated good for another, local labor markets are interconnected via the mobility of workers and their choice in which market to locate. Worker mobility gives rise to population dynamics, and the population dynamics result in stable patterns of the city size distribution that are remarkably constant across countries, known as Zipf's law and Gibrat's law for cities (Eeckhout 2004). Much of the location decisions by workers are governed by prices, in particular wages and housing prices.

My goal here is to review a recent line of research that examines how the location decision differs for workers with heterogeneous skills. I will refer to such skill-dependent location decisions as spatial sorting. Spatial sorting may help explain the urban wage premium, one of the most robust stylized facts in the literature. Wages are higher in larger cities, and this may well be driven by the

disproportionately high presence of high skilled workers. In this *opuscle* I discuss the method for analyzing the skill distribution across cities as proposed in Eeckhout, Pinheiro and Schmidheiny (2014). I will discuss how this method allows me to characterize spatial sorting and provide evidence whether spatial sorting is a determinant of the urban wage premium. To this end, it measures the skill distribution using wages and housing prices.

The ultimate goal is to understand the technological determinants of spatial sorting. Using theory, the variation in the skill distribution across cities can be related to different degrees of complementarities between skilled labor inputs. Finally, the theory also allows to back out the distribution of city-specific productivity across locations, and to study the impact of taxation on the efficient allocation of resources.

2. The Urban Wage Premium

In large cities, workers earn higher wages. This fact has long been established and is known as the Urban Wage Premium. The magnitude of this premium is sizable. The elasticity of wages with respect to city size is 0.042 using US data on metropolitan areas (Eeckhout, et al., 2014). This is illustrated in Figure 1 where each observation corresponds to one city, with its population size in logs on the horizontal axis and the average log wage on the vertical axis. As the population size doubles, the average wage increases by 4.2%, the slope of the regression coefficient depicted in the figure. The difference therefore between small cities such as Janesville, WI with a population of 160,000 and the New York metro area with a population of 19 million is 22%. If the wage reflects the productivity of labor as it does in a compet-

Figure 1. The Urban Wage Premium
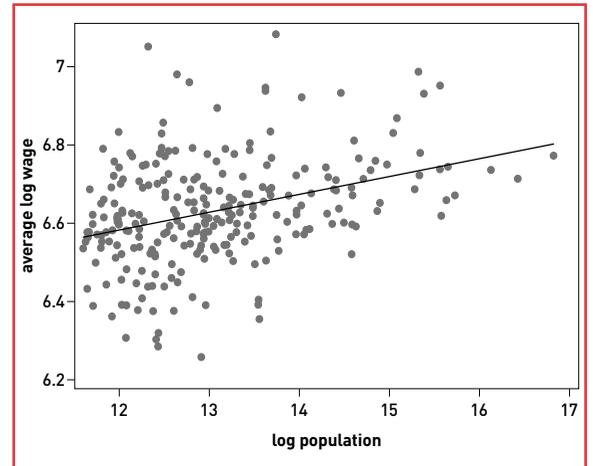
itive market, then this provides evidence at the same time of an urban productivity premium: larger cities are more productive than small cities.

This fact has been known for a long time. The higher productivity of cities has been observed since Adam Smith. Empirically, it was more formally established first in Black (1936) and Schultz (1945) who discuss the urban-rural wage differential during the great depression of the 1930s. Glaeser and Maré (2001) corroborate the finding by providing evidence that the fact is robust across time and space.

Yet, there is an unresolved puzzle. It is not clear what the role is of skills in determining the urban wage premium. Is this wage difference driven by mere productivity differences across cities or by variation in the skill composition, or both? If big cities attract more skilled workers, the wage premium reflects sorting of skills into cities. High skilled workers earn higher wages than low skilled workers and if there are more high skilled

workers in large cities, average wages are higher. Hence the urban wage premium. Instead, if the skill composition across cities is identical, then the wage difference stems directly from the fact that jobs in larger cities are more productive. Identically skilled workers produce more output, which is reflected in higher wages. When the skill composition does not vary across cities, then the urban wage premium is fully explained by productivity differences and not by systematic differences in the skill composition. In what follows, I will refer to any systematic difference in the skill composition across cities as *spatial sorting.*

The decomposition of the productivity differences into sorting of workers and another component due to city-specific shifters builds on a theoretical foundation in which aggregate productivity consists of total factor productivity (TFP) as well as skilled labor input. This is borrowed from the macroeconomic literature that explains differences between countries. While TFP is taken as exogenous, this is only a metaphor for endogenously generated agglomeration externalities that drive these TFP differences, just like the endogenous growth models driving TFP differences across countries. There is an extensive literature analyzing the role of agglomeration externalities that provide appealing explanations for why those differences in TFP arise, ranging from knowledge spillovers between innovation firms, over job search externalities due to labor market density, to consumption externalities and the supply of a variety of goods. For excellent overviews on this topic, see amongst others Duranton and Puga (2004), Rosenthal and Strange (2004), and Moretti (2011). For the remainder, we assume decreasing returns in production at the firm level, with TFP that is city-specific.

In this context I analyze the spatial equilibrium where individual agents take wages and house prices as given and freely choose where to locate. The free mobility assumption will in equilibrium lead to utility equalization across locations for identically skilled workers.

Our objective is two-fold. First, a measure of skills, and hence the skill distribution across different cities, must be derived. Knowledge of the skill distribution provides information whether or not there is spatial sorting, and skills are systematically distributed over cities of different sizes. Once the skill distribution is derived, the impact on wages can be disentangled of skills versus that of city-wide productivity. The analysis builds on Eeckhout, et al. (2014).

3. Measuring the skill distribution

To evaluate the allocation of skills across cities and hence its role in determining the wage premium, the impact of other determinants that affect the location decision of workers must also be taken into account. Equalization of utility across cities for identical individuals implies that differences in real wages can be decomposed into differences in individual productivity —i.e., sorting by skill across cities— and differences in city amenities reflected into compensating differentials.

Since both individual productivity and city amenities are empirically hard to measure, this decomposition implies two avenues for measuring skills. One possibility is to measure skills by their observable determinants such as education and experience, hence attributing residual differences in real wages to amenities rather than unobserved skill. The other possibility is to measure skills by real wages, hence assuming no systematic variation in amenities across cities.

The first and most prominent one is the role of differential amenities. Amenities may be hard to identify and even harder to quantify, but often factors related to weather (hours of sun, temperature,...), geography (views over mountains and the sea) and services offered (opera, theatre,...) are singled out. Amenities can of course also be individual specific, such as the presence of family and friends, or a love for a particular neighborhood. Amenities will only affect aggregate location decisions as long as there is a common preference component. In cities with better amenities, workers are willing to accept lower wages because they are compensated with the value of the amenities. While amenities are real and an important determinant of wage differences between local labor markets, recent evidence (Albouy 2009) shows that there is no systematic relation between amenities and city size (though Glaeser and Gottlieb (2006) do find evidence). For the remainder of this review, I will work under the assumption that amenities do not drive the urban wage premium, but more work is needed to incorporate the role of amenities.

In addition to amenities, the second determinant other than wages that affects the location decision of workers across different cities is the cost of living, measured by housing prices. If those wage differences between large and small cities are so big, why do not all workers move to the big cities? If an economy is to be in equilibrium and workers have no incentive to relocate, there must be a countervailing force that stops workers from moving into the high wage cities. The force pulling in the opposite direction is the cost of housing which is systematically higher in large cities. There are of course also frictions due to the relocation cost, but these cannot explain why people would forego a wage premium as sizable as the one documented above. For one, relative to the permanent life time income of a worker, the cost of moving once is small. Second, gross migration outnumbers net migration by a factor of four. This indicates that people move back and forth between cities much more often than what is needed to equilibrate the market and arbitrage away any differences. Therefore, housing costs are the main countervailing force against the wage differences.

As systematic as the urban wages premium is the fact that house prices increase in city size. Living in big cities is substantially more expensive than in small cities. However, there is a complication. Prices cannot simply be backed out from transaction prices as the quantity of housing consumed varies with prices. Those living in expensive cities facing high housing prices will adjust their consumption of housing by substituting away from housing into consumption goods. In fact, there is consistent evidence that the expenditure share of housing is remarkably constant across different cities. This means that people in large and expensive cities spend about the same share of their income on housing as those in small cities. Because the unit price of housing (per square meter) is so much higher, they must necessarily consume a smaller quantity. Apartments in Manhattan are remarkably smaller than detached houses in Janesville, WI.

To obtain a measure of the unit housing cost Eeckhout, et al. (2014) use data from the American Community Survey by the Census Bureau. They run a hedonic regression of the total housing cost (price times quantity) on a set of observables that correlate with the size of the house such as the number of rooms and bathrooms, whether the house is detached, whether it is an apartment, etc. This gives them an index that they interpret as the measure of the unit cost of housing and they find that the elasticity of housing rental prices with respect to city size is 16.9%. As the population size doubles, the cost of housing increases by 16.9%. As a result, the price of land per square

meter in New York is 147% higher than the price in Janesville, WI. They obtain similar elasticities when using prices for housing in property rather than rental.

The wage data together with the price index then allows them to construct their measure of skills. Based on the preferences for housing and consumption they can derive the indirect utility that depends on wages, house prices and the preference for housing consumption that can be derived from the observed expenditure share of housing. The indirect utility is monotonic in skills, and therefore they can use this price-based measure as a proxy for skills. This includes both observed and unobserved skill heterogeneity.

4. Spatial sorting

I now focus on the spatial sorting of workers by skill and evaluate how the allocation of skills varies across different local labor markets. Ultimately, this will allow to disentangle the impact of the urban wage premium due to productivity differences and that due to sorting of workers. Because there is evidence that capital is fairly efficiently allocated within countries (see Kalemli-Ozcan, et al., 2010), I focus on skills only, and more importantly, on the distribution of skills within cities. Is it that big cities are more productive and have higher measured output because they attract on average more skilled workers?

Based on the method described above that adjusts wages for the cost of living, the distribution of skills can be backed out within cities and how these distributions vary across cities of different sizes. The first finding is that the average of this measure of skills is remarkably constant across cities. In Figure 2, the distribution of skills (or real

Figure 2. Skill distribution for small and large cities, Kernel density estimates
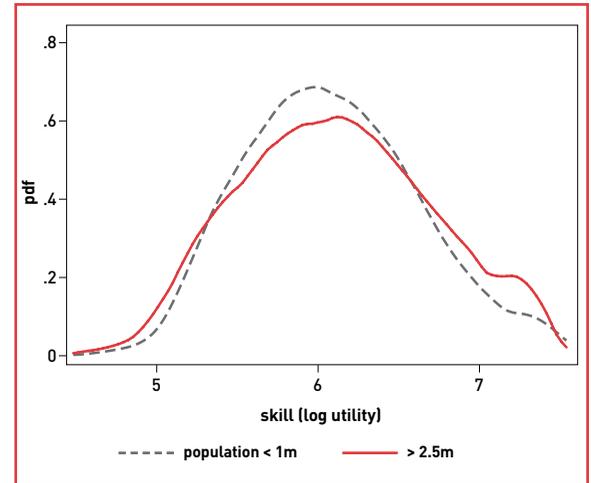
wages) for two groups of cities is plotted. The dashed line depicts the wage distribution of small cities with a population less than one million, and the solid line, the wage distribution of the large cities, with a population larger than 2.5 million.

The distribution of both groups of cities shows a common mean, which is confirmed in more formal tests using individual level observations without grouping the populations. The main implication of this finding is that the urban wage premium is not driven by spatial sorting. On average, the skill level in cities of different sizes is the same, and therefore the observed wage differences are not due to a workforce that is on average more skilled. This is an important finding. A long established fact, the urban wage premium, is exclusively driven by housing price differences, and not by spatial sorting as shown here.

The fact that spatial sorting does not explain the urban wage premium does not mean that there

is no spatial sorting altogether. Quite to the contrary, there is a distinctly systematic pattern of skill allocation across cities. The distribution of skills within cities varies across cities of different sizes exclusively in the standard deviation of the distribution. As illustrated in Figure 2 above, the distribution of skills in large cities is centered around the same means as that in small cities, but it has a larger variance. This finding is very robust and holds true also for example when the skill distribution is analyzed at the individual city level. The scatterplot in Figure 3 establishes that the standard deviation is significantly increasing with city size.

These findings can also be confirmed by looking at observable skills rather than our price-based (wages and housing prices) measure of skills. This is in line with earlier work by Berry and Glaeser (2005), Bacolod, Blum, and Strange (2009), and Combes, et al. (2012) (the last authors find a higher average level of skills when measuring skill by density rather than by population). Other studies, Glaeser and Maré (2001) and Combes, Duranton, and Gobillon (2008), use worker fixed effects to back out unobserved skill rather than using a price-based measure of skill. Baum-Snow (2012) and De la Roca and Puga (2012) have exploited the panel dimension in a dynamic setting to measure skills.

Finally, in recent work, Behrens, Duranton and Robert-Nicoud (2014) analyze the distribution of heterogeneous agents across cities. Their model predicts perfect sorting by talent: New York attracts all the Ph.D.'s, Los Angeles and Chicago all the Masters,…, and all the high school dropouts locate in small cities like Janesville, WI. Within-city heterogeneity in productivity is due to an ex post shock upon 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 the model
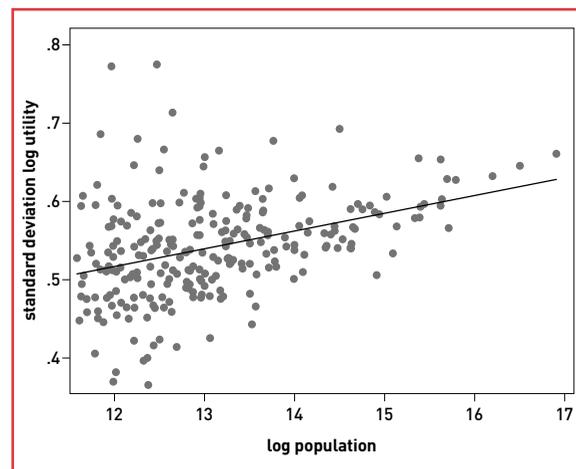
by Eeckhout, et al. (2014) predicts and these last authors also found in the data. The Behrens, et al. (2014) result of perfect sorting in talent is a direct consequence of assuming that each worker consumes one unit of land independent of his wage. Highly talented workers with high wages are therefore, relatively, much less affected by high housing costs in large cities than less talented workers. In contrast, Eeckhout, et al. (2014) do allow housing consumption to increase with wages and hence with talent in line with overwhelming empirical evidence. In the equilibrium allocation of the 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. Most importantly, it also allows us to infer the pattern of complementarities in the technology that drive the location decision of workers. I turn to this next.

5. An explanation: skill complementarity

An important question is how this pattern of sorting can be explained. Sorting patterns in general are driven by complementarities between inputs. For example, sorting of workers across firms or teams are driven by the extent to which workers of different skills are complementary with each other. The stronger the complementarity between given types, the higher the productivity and the wages, and therefore the higher the fraction of those workers in that firm. Extending this logic to sorting between cities, observing a skill distribution that has a constant mean across cities but higher variance and therefore thicker tails informs us about the underlying technology. A higher fraction of both the low skilled together with high skilled, all the while with a lower fraction of medium skilled implies that there are stronger complementarities between the high and the low skilled than between the high and the medium skilled.

This pattern of complementarity can be interpreted as a source of demand for low skilled services by the high skilled. Highly skilled professionals tend to congregate in large cities because the productivity is highest there. The high productivity in large cities is complementary with skills and generates higher output and thus higher wages. This in turn leads to higher demand for those complementary skills. Highly specialized surgeons, for example, who join the Memorial Sloan-Kettering Cancer Center in New York will be surrounded by medium skilled nurses and collaborators. Yet, they demand even more low skilled services such as administrative help and cleaning services at work, and at home the low skilled services from day care, restaurants and maids. The relative higher demand for low and high skilled work that is observed in the skill distribution in large cities can thus be explained by the relative stronger complementarity between high and low skilled inputs. The high opportunity cost of time for the high skilled leads to a larger demand for the low skilled services.

Key in the explanation based on extreme skill complementarities to explain the thick tails in the observed skill distribution is the complementarity between TFP —whether it is exogenous or endogenous due to agglomeration externalities— and the skill aggregator. To see this, observe that also in small cities there is extreme-skill complementarity. However, because of the TFP-skill complementarity, the effect of the between-skill complementarity is accentuated in cities with larger TFP, and hence in larger cities.

In an attempt to explain this pattern, Eeckhout, et al. (2014) construct a model of city choice based on wages and cost of living. The key feature is perfect mobility of citizens with different skills. This induces equalization of utility across different locations for equal skills. Wages and house prices are determined in general equilibrium together with this arbitrage condition from mobility. They show that a technology of production at the firm level with heterogeneous skill inputs generates the observed pattern of spatial sorting when the elasticity of substitution systematically varies across skills. In fact, as a benchmark in this model, they find that a technology with a standard constant elasticity of substitution across skills necessarily leads to identical skill distributions across cities of different sizes.

This has also important implications for our understanding of production technologies that have constant elasticities of substitution. While analytically very tractable and therefore desirable tools for modeling, constant elasticity of substitution technologies cannot match many features of

the equilibrium allocation of skills, in our setting to cities, but more generally. In fact, the technology that I have discussed above, the nested constant elasticity of substitution production function, has a counterpart in other work. Krussell, et al. (2000), for example, use a technology with nested elasticities of substitution to explain the evolution of the skill premium. They show that a technology with stronger complementarity between high skilled labor and capital generates an increasing wage premium over time. As output grows, capital investment increases and as a result of the stronger complementarity, the marginal product of high skilled labor increases with increasing capital investment.

Instead of an intertemporal effect of productivity on inequality, I find a similar effect across space. This cross-sectional effect can be understood as follows. As the productivity of a city is larger, the marginal product of both high and low skilled labor increases, leading to larger demand for these extreme skills. This in turn translates in bigger inequality in those more productive cities, and those turn out also to be the larger cities.

This measure of skills is completely price-based. In fact, it only uses information on wages and housing prices and no other observable characteristic. There is a good reason for doing so. Only about one third of the wage variation can be explained by observable characteristics such as education, occupation, industry, and personal characteristics. Observable characteristics are at best only an incomplete measure of skills. The lion's share is determined by unobservable characteristics such as non-cognitive skills. For example, the ability to interact socially, or to arrive on time and to meet deadlines.

Nonetheless, it is of interest to investigate whether the distribution of observable skill charac-

teristics exhibits the same pattern of spatial sorting where the average skill is constant across city sizes and the standard deviation is increasing in city size. Separately conditioning on years of schooling and on occupational indices, I find that the same results continue to hold. The skill distribution has a constant mean and an increasing standard deviation. I also investigate the role of industries for different cities. So far, I have started from the premise that different cities have different TFP but the same production technology. If instead different industries have different technologies —-say banking in New York versus auto manufacturing in Detroit— then thick tails may be driven by industry characteristics. It turns out that the thick tails result continues to hold when conditioning on industry codes. Finally, personal characteristics could help explain the thick tails result. Migration for example, can drive foreign-born workers into big cities. If those migrants are predominantly low skilled, then there would be thick lower tails simply due to migration. Splitting the sample into foreign-born and native workers, however, continues to generate thick tails for both samples. Likewise, there may be a life cycle effect due to learning. Young, low skilled workers may go to the big city to try their luck and as they grow older move back in case they do not make it, and stay if they do. This would generate a lower and an upper thick tail as those returning would have accumulated some skills without becoming high skilled. Splitting the sample by cohort, however, continues to generate the thick tails for each of the cohorts, thus corroborating the robustness of the result.

To illustrate that unobservable characteristics are the dominant force also in driving the tail difference between large and small cities, Eeckhout, et al. (2014) perform a decomposition exercise of the quantiles of the distribution. The aim is to explain how much of the difference between the tails in large and small cities is explained by observables,

including all the ones mentioned above as well as gender and race. The unobserved characteristics explain the majority of the difference in the tails, but there is an asymmetry. In the upper tail, the observed characteristics explain just under 50% of the variation whereas in the bottom tail, observables explain only one quarter. This further confirms that using a price-based explanation for the measure of skills is justified. But also, it suggests that unobserved heterogeneity as a source of this tail difference is even more important for the low skilled than for the highly skilled. This may suggest that non-cognitive skills are even more important for the low than for the high skilled workforce.

I can thus summarize the role of spatial sorting as driven by differential degrees of relative complementarity. This leads to thick tails in the skill distribution but no difference in the means. As a result, while spatial sorting is an important consequence of technological complementarities, it cannot explain the urban wage premium. The reason why higher wages can be sustained in larger cities is because they are offset by higher housing prices.

6. Productivity differences across cities

The process of evaluating the role of skills in the determination of wages also generates information on the variation of Total Factor Productivity (TFP) across local markets. As I pointed out above, in the presence of competitive labor markets, wages reflect the worker's productivity which is determined by the quality of the input —the worker's skill— and the location specific TFP.

In a competitive labor market, wages reflect the marginal product of workers. Wages are higher in large cities because the value of their output is higher. This can be due to agglomeration externalities as first pointed out by Marshall (1890). With more workers close together, new technologies spread faster and more efficiently, hiring is facilitated and trade is more interactive and specialized. Adjusting wages for the cost of living, we can have an idea whether there is a systematic difference across cities. When there is sufficient mobility, identically skilled workers will be indifferent between living in a large city and a small city. Therefore, they must choose bundles of wages and housing prices that generate the same utility. With the expenditure share on housing constant, utility equalization across different size cities implies that wages adjusted for housing prices, or real wages, must be constant. Because those real wages reflect utility, they must necessarily also reflect skills if workers are heterogeneous.

Moreover, there is ample evidence of labor mobility of workers across the US. In the late nineties for example, in response to higher wages in the Bay area due to the technology shock from the tech boom, there was a major influx of labor. The inflow of labor stops and the economy is in equilibrium once housing prices adjust sufficiently to counter the benefits from higher wages.

Based on the model specification with representative workers, and using data both on wages and the price of housing (adjusted for the quality of housing), I can derive the distribution of TFP across locations. For the US data, the model can be used to back out the value of TFP for different cities. For each city, average wages and the population size are observed. Since wages reflect marginal product, given the production technology this allows to back out the TFP for each city.

Population mobility implies equalization of utility across cities for identical agents. Hence, any remaining variation in utility that leads to devia-

tions from utility equalization for observed populations are attributed to amenities. Even though the role for amenities is disregarded in systematically explaining differences by city size, amenities are important in generating utility equalization for observed wages. The largest cities are indeed the most productive in general. It turns out that there is substantial heterogeneity across locations, with the highest productivity locations having a local TFP that is three times as high as that in the lowest productive cities. In Table 1 I report the estimated values for TFP (denoted by A, normalized to one for the country average) and for the value of amenities (denoted by ε, normalized to zero for the country average) for those cities at the extremes of the TFP and amenities distribution.

The most productive cities are Bridgeport-Stamford-Norwalk, CT where many of the hedge funds are located, as well as the main urban areas in Silicon Valley. The least productive metro areas are small towns in Texas and Bowling green. The average productivity differences are astonishing: an average worker in Silicon Valley produces three times as much output as the average worker in Bowling Green, KY, taking into account that this is the average productivity for a distribution of skills that is invariant in the mean.

There are also substantial differences in amenities. It happens to be the case that the largest cities in the sample, New York, Los Angeles and Chicago also have the highest amenity levels. However, this relation between city size and amenities is not borne out over the entire sample. In fact, the correlation between TFP and amenities is zero. Maybe there is something especially valuable in terms of amenities of the extremely large metro areas. The cities with lower amenity values are Saginaw, MI, Athens, GA, and Ocean City, NJ. The latter indicates that the presence of the coast does not guarantee a high amenity value.

The measurement is crucial for the aggregation of TFP. Consider the parallel with unemployment. When unemployment is not properly accounted for, then with high levels of unemployment, productivity will appear higher either because the basis over which productivity is calculated is wrongly measured, or because the marginal productivity of labor varies with the quantity of labor. Similarly, when aggregating TFP across locations, one needs to properly take into account the underlying inputs, and skills in particular.

Despite the huge differences in output and productivity, mobility between locations implies utility equalization and therefore no utility gain from higher productivity in large cities. Living in large cities requires paying higher housing prices and living in smaller dwellings, to the point where all the wage and productivity gains are eaten away.

Table 1. Estimated total factor productivity A and amenities ε for selected cities

MSA	A	ε
Highest A		
Bridgeport-Stamford-Norwalk, CT	1.38	-0.16
San Jose-Sunnyvale-Santa Clara, CA	1.36	0.14
San Francisco-Oakland-Fremont, CA	1.35	0.44
Lowest A		
Brownsville-Harlingen, TX	0.53	0.00
Amarillo, TX	0.49	-0.02
Bowling Green, KY	0.46	-0.26
Highest ε		
New York-Northern New Jersey-Long Island	1.17	1.45
Los Angeles-Long Beach-Santa Ana, CA	1.02	1.37
Chicago-Naperville-Joliet, IL-IN-WI	1.06	1.07
Lowest ε		
Saginaw-Saginaw Township North, MI	1.17	-0.46
Athens-Clark County, GA	1.04	-0.53
Ocean City, NJ	1.12	-0.63

The objective therefore is to evaluate the impact of productivity differences from a welfare viewpoint.

Finally, the role of mobility is crucial for the long term determination of the city size distribution. In reality, total factor productivity is not constant. Due to technological progress, the TFP of certain cities evolves. In the Bay area for example, TFP received a series of positive shocks during the tech boom, whereas cities like Detroit or Saint Louis have suffered negative shocks due to the evolution of the manufacturing industry. In a dynamic economy, those random productivity shocks translate into mobility by workers. Surely, wages in Detroit have fallen and workers have been attracted to move to San Francisco where wages have gone up. This is exactly the mechanism that I have laid out above.

7. Policy-induced misallocation

Even if TFP is unaltered and policy cannot change actual TFP, the full potential of TFP to translate the allocation of productive inputs into output may not be fully realized. This may be due to the misallocation of the productive factors. Productive resources may be distorted across cities, thus leading to differences in the output produced. While most of the existing work on cross-country differences looks at capital, here the focus is on the misallocation of skilled labor. In particular, I discuss the impact of federal income taxation on the location decision across labor markets within a country, and the possible distortions in this location decision (based on Eeckhout and Guner, 2014).

Now why does income taxation affect the allocation of resources across cities? Consider the comparison I made earlier between New York and

Janesville, WI. I know that the average skill of a worker is statistically the same in both cities. Yet, wages in New York are more than 20% higher. Given progressive taxation, higher nominal income earners pay higher average taxes. As a result, the representative median worker in New York pays an average of 26.5% in Federal taxes compared to 23.5% in Janesville, WI. Therefore, identical agents pay different tax rates. We can abstract from local taxes since those tend to stay within the city and any tax differences tend to be compensated with differences in services and amenities.

The existence of progressive taxation is taken as given. All developed economies have progressive tax systems where those earning higher incomes pay a higher average tax rate. Amongst the justifications proposed for progressive taxes are efficiency motives such as providing optimal incentives to work, insurance against income risk, consumption smoothing, or equity motives that demand redistribution from the rich to the poor. In this work, the question is how such a tax schedule should be optimally designed to account for the fact that wages for identical agents differ across different locations. The basic underlying principle here is that identical agents should be treated identically by the tax system. Currently, that is not the case since the same skilled worker in Janesville, WI pays a lower average tax rate than her counterpart in New York city.

This has profound implications for optimal taxation. Income taxes do not typically condition on location, and as a result, progressive income taxes levy a higher burden on otherwise identical workers in large cities than on their equally skilled counterparts in small cities. Because there is mobility, those identically skilled workers who are differentially taxed nonetheless still receive the same utility: workers move until they are indifferent between locations. As a result, such progres-

sive taxation does not have any redistributive impact. Of course, redistribution is possible between differently skilled agents. Therefore, the optimal tax schedule is investigated that corrects for the first inefficiency between identical types, while maintaining the progressiveness between different types.

The role of geographical differences and its impact of federal policies has long been recognized by economists such as Wildasin (1980). A recent quantitative assessment of distortions of taxation and their impact on the spatial misallocation of workers is provided by Albouy (2009). Similar points have been made concerning federal transfers (rather than taxes) and the distortions this imposes on people's location decisions, both in the economics as well as in the law literature (Kaplow, 1995; Glaeser, 1998; Knoll and Griffith, 2003). This work points out the distortion of such policies. An open question is what the optimal tax schedule is. An optimal tax schedule will take into account equilibrium behavior by individuals and their location decision affects prices, notably wages and housing prices.

Now suppose a tax schedule is introduced that is location specific, either by conditioning on the city size or by taking into account the cost of living. Notice that the objective of this tax schedule is to impose the same tax burden on equally skilled workers in different cities. This schedule can be progressive in the sense that high skilled workers pay higher average tax rates than low skilled workers. So let the representative agent in New York and in Janesville pay the same tax (say 25%). Then New York becomes more attractive to live and Janesville less attractive. In fact, there will be a net flow of workers from small cities into larger cities. Because large cities have higher TFP, this implies that the productivity of those workers increases. Even the productivity of the workers in

small cities may increase if there are decreasing returns in production. Then the marginal product of workers in small cities goes up at constant TFP if the population declines. But also the cost of living in the large cities increases, and equilibrium is restored if wages adjusted for the cost of living is again equalized across locations. The new equilibrium will have higher output, as each worker is more productive.

There is huge potential for growth, even if TFP itself does not increase. Country-wide output increases can be sizable once the optimal tax schedule is introduced. Productivity goes up for all skills, in all regions and in that sense the policy is Pareto improving. At the same time, to generate those output gains, the policy implies huge changes in the population distribution. And here there is a trade-off. A higher concentration of people in highly productive (and large) cities increases output, but it also increases the cost of living. As a result, a taxation system that takes into account this trade-off will not be flat. The optimal tax will be city specific and will feature a degree of progressiveness that will maximize utility. Also, policy should not be blind to the short term adjustment costs that such mobility entails. And it is paramount to take into account the impact on the distribution of income across heterogeneously skilled workers. The key measure of inequality must be utility (wage income adjusted for housing prices), and not just wages.

For all these reasons, the optimal tax schedule will be progressive even between identically skilled workers. The planner cares about utility, and the worker who earns high wages and generates a lot of output can nonetheless be miserable living in a dense and expensive city. The optimal allocation trades of the productivity gain against the loss in utility from living in small housing and in congested cities.

Conclusion

Local labor markets are substantially differentiated: large cities command higher wages, have higher productivity and workers pay higher housing prices. Yet, those markets are connected through the mobility of the work force and the choice of workers where to locate. The interconnection of local labor markets permits the identification of technological characteristics that determine productivity from the observation of the population distribution over cities.

I have reviewed the results that identify the role of spatial sorting, i.e., the extent to which the composition of skills differs across cities of different sizes. The main finding is that spatial sorting does not explain the skill premium, i.e., the fact that average wages and average productivity is higher in large cities. Cities of all sizes attract on average workers of the same skill level. The wage premium is entirely counteracted by the higher housing prices. Hence, on average there is no higher skill composition.

However, spatial sorting is an important ingredient for the selective skill composition driven by the technology. Large cities disproportionately attract both more high and low skilled workers. This is due to the complementarities between those extreme skilled workers. In combination with higher levels of productivity, those complementarities lead to higher demand for both high and low skilled workers, which translates in the observed thick tails of the skill distribution in large cities. This thus provides evidence of the underlying technology.

I also reported the findings of sizable differences in the productivity across cities. Larger cities have larger estimated levels of productivity. The most productive cities are about three times more productive than the least productive. The former tend to be larger, the latter are smaller.

Understanding the workings of local labor markets also has important policy implications. Federal income taxation affects wages without taking into account the cost of housing. As a result, more productive cities that pay higher wages will lead to higher marginal tax rates when taxation is progressive. Optimal spatial taxation demands a tax schedule that is contingent on the city size and can have substantial impacts on the city size distribution.

References

Albouy, D., (2008), ''Are big cities bad places to live? Estimating quality of life across metropolitan areas,'' NBER Working Papers 14472, National Bureau of Economic Research, Inc.

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

Baum-Snow, N. and R. Pavan, (2012), "Understanding the city size wage gap," *Review of Economic Studies* 79(1), 88–127.

Black, J. D., (1936), "Agricultural wage relationships historical changes," *Review of Economics and Statistics*, 18(1), 8–15.

Behrens, K., G. Duranton and F. Robert-Nicoud, (2014), "Productive cities: sorting, selection and agglomeration", *Journal of Political Economy*, 122(3), 507–553.

Combes, P.-P., G. Duranton, and L. Gobillon, (2008), "Spatial wage disparities: sorting matters!," *Journal of Urban Economics*, 63(2), 723–742.

Combes, P.-P., G. Duranton, L. Gobillon, and S. Roux, (2012), "Sorting and local wage and skill distributions in France," *Regional Science and Urban Economics*, 42(6), 913–930.

De la Roca, J., and D. Puga, (2012), "Learning by working in big cities", CEMFI mimeo.

Duranton, G. and D. Puga, (2004), "Micro-foundations of urban agglomeration economies,", in J. V. Henderson and J. F. Thisse (eds.), *Handbook of Regional and Urban Economics*, edition 1, volume 4, chapter 48, 2063–2117.

Eeckhout, J., (2004), "Gibrat's law for (All) cities." *American Economic Review*, 94(5), 1429–1451.

Eeckhout, J., R. Pinheiro, and K. Schmidheiny, (2014), "Spatial sorting", *Journal of Political Economy*, 122(3), 554–620.

Eeckhout, J., and N. Guner, (2014), "Optimal spatial taxation," mimeo.

Gabaix, X., (1999), "Zipf's law for cities: An explanation." *Quarterly Journal of Economics*, 114(3), 739–67.

Gibrat, R., (1931), *Les inégalités économiques; applications: aux inégalités des richesses, à la concentration des entreprises, aux populations des villes, aux statistiques des familles, etc., d'une loi nouvelle, la loi de l'effet proportionnel.* Paris: Librairie du Recueil Sirey.

Glaeser, E. L., (1998), "Should transfer payments be indexed to local price levels?," *Regional Science and Urban Economics*, 28(1), 1–20.

Glaeser, E. L. and J. D. Gottlieb, (2006), "Urban resurgence and the consumer city," Harvard Institute of Economic Research Working Papers 2109.

Glaeser, E.L., and D.C.Maré, (2001), "Cities and skills", *Journal of Labor Economics*, 19(2), 316–342.

Kalemli-Ozcan, S., A. Reshef, B. E. Sorensen, and O. Yosha, (2010), ''Why does capital flow to rich states?,'' *Review of Economics and Statistics* 92(4), 769–783.

Kaplow, L., (1995), "Regional cost-of-living adjustments in tax/transfer schemes," *Tax Law Review*, 51, 175–198.

Knoll, M. S., and T. D. Griffith, (2003), "Taxing sunny days: adjusting taxes for regional living costs and amenities," *Harvard Law Review*, 116, 987–1023.

Krusell, P., L. Ohanian, J.-V. Rios-Rull and G. Violante, (2000), "Capital-skill complementarity and inequality: A macroeconomic analysis," *Econometrica* 68(5), 1029–1053.

Marshall, A., (1890), *Principles of Economics*, Macmillan, Cambridge, MA.

Moretti, E., (2011), "Local labor markets," *Handbook of Labor Economics*, Elsevier.

Rosenthal, S. S. and W. C. Strange, (2004), "Evidence on the nature and sources of agglomeration economies," in J. V. Henderson and J. F. Thisse (eds.), *Handbook of Regional and Urban Economics*, edition 1, volume 4, chapter 49, 2119–2171 Elsevier.

Schultz, W., (1945), *Agriculture in an Unstable Economy*, McGraw-Hill Book Company, Inc., New York.

Wildasin, D. E., (1980), "Locational efficiency in a federal system," *Regional Science and Urban Economics*, 10(4), 453–471.

Opuscles already published

Jan Eeckhout

Jan Eeckhout is Professor of Economics at UPF, Barcelona GSE and UCL. He has teaching and research interests in applied economics, with a special emphasis on the labor market. He studies unemployment, the determinants of compensation, and inequality in cities. His work has been published in the *American Economic Review, Econometrica*, the *Review of Economic Studies*, the *Journal of Political Economy*, and has been supported by several government grants, including funding from the National Science Foundation (US) and the European Research Council. Jan Eeckhout has been a tenured professor at the University of Pennsylvania, where he was for 9 years. He has also taught for one year at NYU Stern and visited MIT. He has been an editor of the *International Economic Review* and is on the editorial board of the *Review of Economic Dynamics*, the *Journal of Economic Theory* and the *Journal of the European Economic Association*, and he is a fellow of the European Economic Association. He received his Ph.D. in Economics from the London School of Economics in 1998.

CREi

Centre de Recerca
en Economia Internacional

Academic unit of the
barcelona|**gse**
graduate school of economics

**Universitat
Pompeu Fabra**
Barcelona

**Generalitat
de Catalunya**