The Global Value of Cities

Aakash Bhalothia, UC San Diego
Gavin Engelstad, Macalester College
Gaurav Khanna, UC San Diego*
Harrison Mitchell, UC San Diego

January, 2025
Preliminary

Abstract

We estimate the economic value of each city around the world, and how it changes over the process of economic development. We obtain a dataset of detailed job histories of 513 million workers in 220,000 cities across 191 countries. We track job spells for these workers, with information on start and end dates, establishment name, location, job title, and effective salary. We use an event-study movers design, conditional on individual and time fixed effects, with tests for pre-trends. We show that moving to a higher (average) wage city leads to a substantial jump in earnings. The distinctly global nature of our analysis allows us to measure internal and cross-border moves, while simultaneously estimating what makes a city more productive, and how the gains from spatial re-sorting of workers vary over the trajectory of economic development. When moving across international borders, 93% of the change in wages can be attributed to city effects, whereas within countries about 45-73% of the wage change is because of city effects. Richer countries exhibit stronger ability-based sorting, reducing the proportion of wage differentials explained by city effects. City effects are strongly linked to economic structure, whereby cities with greater industrial diversity exhibit larger city effects, and more populous cities tend to be more productive, consistent with agglomeration economies. The variance in the contribution of city effects within countries highlights the potential gains from migration. These gains are particularly large in low-income, low-urbanization countries, underscoring the importance of reducing migration frictions in the developing world.

Keywords: City wage premia, movers design, spatial sorting, wage differentials

*Corresponding author: gakhanna@ucsd.edu

We thank Krystal Qiu and Akshay Yeddanapudi for excellent research assistance. We thank Victor Couture for generously sharing his data. We are grateful to various seminar participants and colleagues for their insightful comments.

1 Introduction

What are the income gains from moving a software engineer from Kansas City to San Francisco? Or from Bangalore to San Francisco? Estimating these gains is challenging, as observed wage differences across cities may primarily reflect the sorting of individuals based on skill rather than the causal effect of location. To isolate the true value of a city on worker's earnings, we would need to track cross-city moves for a large number of workers. We leverage a unique new database of high-skilled worker job histories for 513 million workers in 220,000 cities around the world. With the help of an event-study movers design that exploits worker movement (Finkelstein et al., 2016; de La Roca and Puga, 2017; Card et al., 2023) we ask: are the large disparities in city wages across the globe primarily driven by sorting on skill (Young, 2013; Behrens et al., 2014; Combes et al., 2008), or do cities themselves have a meaningful impact on a worker's earnings (Ciccone and Hall, 1996; Diamond, 2016; Glaeser and Gottlieb, 2009; Card et al., 2023)?[1]

In this paper, we distinguish between *city effects* and *individual-level sorting*. City effects capture a city's productivity, its firm activity, the types of jobs, technology, infrastructure, and agglomeration forces it may have. In contrast, systematic sorting occurs when high-ability individuals self-select into more desirable locations, leading to wage differences that primarily reflect worker characteristics rather than city-level productivity. Distinguishing between these explanations has meaningful implications for how facilitating the movement of people can affect aggregate income. If city effects are indeed important, then migration from low-productive to high-productive cities can raise aggregate productivity. If, instead, sorting on abilities is the primary driver of wage differentials, then cities may instead focus on attracting high-ability workers.

We first estimate city effects for 123,431 cities around the world. We estimate event-study designs based on those who migrate between cities, and quantify what fraction of the wage differentials between any two pairs of cities translate to higher wages for those who move between these cities. Next, we estimate the city effects, accounting for the heterogeneity in wage premia across firms within a city. We then investigate what city features are associated with higher city wage premia. How much of the city effects are a reflection of the nature of economic activity, the types of firms, and industries? We seek to unpack how a city's economic structure is associated with the city effects. Finally, we explore how the potential gains from spatial re-allocation within a country varies across the economic development process. The dispersion of wage premia across cities, within a country reflects how greater internal migration may potentially lead to higher aggregate incomes via workers sorting to high-wage cities. We study how this potential gain from spatial sorting is associated with a country's geographic and economic structure.

We obtain global data on 700 million LinkedIn users that track the movements of 513 million workers across 220,000 cities around the world. Our data include job location, job titles, firm/establishment names, and the start and end date of each work spell. We construct individual-

[1] Pritchett (2017) argues "Mostly in the world there aren't poor people. There are people in poor places."

level panels using the full job histories of individuals. LinkedIn data capture relatively high-skilled individuals. Complementary data on 200 million salaries across a wide variety of sources, by precise location, job title, tenure, seniority, and company (establishment), allow us to track predicted wages, adjusted for purchasing power parity, for individuals. An important feature of our data is that we observe precise granular location information, regardless of where a firm is headquartered. This allows us to track the movements of individuals across cities across the world, and their corresponding changes in job status, titles, seniority, type of work, and effective salaries. We use this information to understand how moving to a particular city affects individual outcomes.

We first highlight descriptive evidence with an event-study analysis of changes in individual wages associated with bilateral moves across any two pairs of cities.[2] Conditional on individual, time, and time-since-move fixed effects, we study how moving to a particular higher-wage city from a particular low-wage city affects an individual's earnings. This design has a few distinct advantages. First, by conditioning on individual fixed effects, we account for all time-invariant characteristics of each worker. Second, the event-study framework enables us to assess pre-trends in wages before a move, helping to determine whether migration was driven by positive or negative wage shocks. Finally, our identification strategy relies on the assumption that wage shocks occurring at the time of relocation are uncorrelated with individual-by-city match-specific factors. If this assumption were systematically violated, we would expect to observe pre-trends in our event-study estimates.

We then exploit the variation across movers, to estimate how moving across cities affects wages. The advantage of individual-level fixed effects, and time-varying city fixed effects is that the estimation allows for systematic mobility related to individual and city characteristics. For instance, more productive workers may be more mobile, more productive cities may be in higher demand, and assortative matching may exist (more productive workers move to high productive cities). Yet, we still need to rely on standard assumptions in the literature on exogenous mobility, and additive separability in these effects. Exogenous mobility requires that mobility is not driven by city-individual match quality, and that shocks (or time-varying factors) do not drive the moves. That is, large job losses or human capital accumulation would violate these assumptions. Our empirical tests, and event studies show a lack of meaningful pre-trends, providing support for these assumptions. We further show that there is symmetry in the effects of moving between pairs of cities, which support assumptions around match-quality driven moves, and additive separability.

Yet, the simple Abowd et al. (1999) AKM framework may not provide us with interpretable estimates if there is heterogeneity in the premiums paid by different firms within a city (Card et al., 2023). If workers who move from a low-wage city to a high-wage city tend to come from firms that paid well in their origin city but moved to firms that paid below-average in their destination city, our estimated city effects would not capture the effect of randomly moving a worker between any two pairs of cities. So, in order to estimate city effects, we estimate 'establishment-level effects,'

[2]The closest in methodology is Finkelstein et al. (2016), who study health outcomes as older Medicare individuals move across the US.

for the universe of establishments in our data, and aggregate them up to the city level. We follow this as our main strategy to measure city effects.

Our analysis highlights six facts. First, there is stark regional heterogeneity in average city-level wages across the world. These reflect regional incomes (city-level wages are lower in Africa and South Asia, in comparison to Europe), country borders (there are sharp jumps at the US-Mexico border), and also within-city regions (coastal USA has higher-wage cities than the midwest).

Second, we examine the sample of cross-city movers (as they help estimate city effects), and find that most moves are concentrated to relatively higher-wage cities. That is, individuals are far more likely to move to a relatively richer city than a relatively poorer one. When looking at cross-country moves, individuals are far more likely to move to a city in the top income quintile of cities in the destination country.

Third, we study bilateral moves between any two pairs of cities, and estimate event-study specifications at the individual-move level to see how wages for workers change for different kinds of moves. The event studies do not suggest any pre-trends, and are robust to various controls, but show a substantial increase in earnings when moving to a relatively high-wage city. These effects are symmetric for moves in the opposite direction.

Fourth, we estimate city-level productivities, and find substantial heterogeneity in the share of geographic wage variation explained by these city effects. Specifically, we find that city effects account for a large share of variation in wages among cities across the world, but explain less within-country variation. We show that for international moves, 93% of the wage differentials between any two pairs of cities, translate into gains for the movers. Whereas, for within-country moves, about 45-73% of the wage differentials translate into gains.[3] As such, country effects are important, but city effects are far from negligible. Additionally, the extent of within-country variation explained by city effects varies across countries. This has important implications for potential gains from migration in different contexts.

Fifth, we correlate these isolated city effects, and find some stark relationships. City-level diversity and complexity in industrial composition are strongly correlated with larger city effects. In contrast, cities with few concentrated industries have lower city effects. Further, larger cities have higher city effects, suggesting meaningful agglomeration economies. We also examine the types of jobs a city has, and find that cities that have more skilled and senior-level jobs have higher city effects. These patterns not only validate our methods for estimating city effects, but also help answer an age-old question in urban economics: what makes a city more productive?

Sixth, we study how the potential gain from spatial sorting varies across the development process. We estimate the variance in the importance of city effects within a country. A larger variance suggests more potential gains from internal migration. That is, large city effect differentials imply that moving individuals to high-impact cities can raise aggregate productivity. We find, in

[3] 50.8% of wage differences in the US reflect city productivity, almost precisely mirroring the 50% finding from Card et al. (2023) on US commuting zones, using the LEHD data.

fact, that the dispersion in the importance of city effects (and gains from internal migration) is larger in low-income countries, with fewer large cities and less urbanization. Indeed, the share of wages explained by city effects is higher in countries with lower migration intensity, suggesting that migration frictions prevent efficient sorting across cities, and so raise the importance of city effects.

These findings have important implications for recent debates on the value of locations. A recent starting point of our work is Card et al. (2023), who use AKM methods to isolate gains from locations within the US. The distinctly global nature of our analysis creates certain advantages. It allows us to not only study city effects within countries across the world, but also cross-country and cross-region movements. These moves reflect important nuances. First, cross-border moves have much larger wage gains. This speaks to important work on the gains from migration, much of which has focused on internal migration within one country (Bryan and Morten, 2019; de La Roca and Puga, 2017; Dauth et al., 2022; Card et al., 2023) or country-level gains (Amanzadeh et al., 2024). Second, our global analysis allows us to understand how the potential gains from reallocating workers across space vary greatly across countries around the world. This reflects recent work on how gains from reallocating inputs in firms and sectors (Hsieh and Klenow, 2009; Gollin et al., 2014). The methods we use are different from earlier work (with some exceptions, like Finkelstein et al. (2016)), as our event-study design allows us to cleanly test for pre-trends that would be indicative of violations of the identification assumptions.

Next, we unpack the correlates of these city effects, to hint at work on why some cities are more productive than others. Previous work highlights the importance of agglomeration economies (Ciccone and Hall, 1996; Duranton and Puga, 2004; Rosenthal and Strange, 2004) and the skill-composition of the workforce (Glaeser and Maré, 2001; Moretti, 2004; Diamond, 2016; Eeckhout et al., 2014) within particular countries. We show global evidence of these patterns, and how they change across the process of development. For instance, agglomeration forces and skill concentration seem to be more relevant in developing countries.

Finally, we speak to the debate on whether the disparities in city wages reflect ability-biased sorting (Young, 2013; Behrens et al., 2014; Combes et al., 2008), or whether locations have a causal impact on earnings (Ciccone and Hall, 1996; Diamond, 2016; Glaeser and Gottlieb, 2009; Card et al., 2023). We do find a meaningful role for ability-based sorting, whereby city effects can sometimes only explain about 45% of the wage differentials across cities within a country. Yet, we argue that these are non-negligible city effects, and are correlated meaningful measures of economic structure.

Our paper is organized as follows. Section 2 describes the data, and how we construct our estimating sample. Section 3 provides some descriptive facts about the distribution of wages in cities across the world, on migration patterns between different types of cities, and our event study designs. Section 4 describes our empirical strategy to estimate city effects. Section 5 discusses the estimated city effects and their implications for heterogeneity in gains from different types of migration. Section 6 unpacks what explains these city effects, and the distribution of city effects across countries. Section 7 concludes.

2 Data

The data for our primary analysis comes from Revelio Labs, which provides detailed data from public professional networking websites, including LinkedIn. While our data contains harmonized individual-level data from other sites in addition to LinkedIn, we will focus our discussion on LinkedIn data since that is a large majority of our users. LinkedIn contains over 700 million professional profiles, which are created by users by entering their personal education and employment histories. We primarily work with the employment information users provide, which contains the location of the job, job titles, firm names, and the start and end date of each entry. By leveraging the start and end date of each position, we are able to construct user-level panels that allow us to follow the same user across the entirety of their work history. Our data from Revelio Labs covers over 500 million accounts from 180 countries as of late 2023.

The type of individuals with LinkedIn profiles are different than the general population. While LinkedIn has profiles with extremely varied histories, highly educated individuals and those with white-collar jobs are overrepresented in our data. We, therefore, view our analysis as informative of the dynamics of earnings across space for the relatively high-skilled.

Additionally, countries have varying degrees of LinkedIn use, even among the college-educated individuals in the country. As noted in Amanzadeh et al. (2024), the ratio of college-educated individuals to LinkedIn users across countries is heterogeneous, with countries in Central and East Asia particularly unlikely to use LinkedIn. Although this could complicate matters when analyzing certain questions, most of our analysis will focus on estimating the city wage premia within a single country at a time.

In order to measure the salary of individuals, we make use of supplementary imputed wage data from Revelio Labs. We observe an individual's imputed salary for each job position on their profile. Wages are estimated using job titles broken down into 1500 categories, company-specific information, geographical, economic information such as median housing values and unemployment rates,[4] position-specific information such as tenure and seniority, and company identifiers. The documentation provided by Revelio explains that wages are imputed using a regression-based model on a dataset of over 200 million salaries across a wide variety of sources. In order to provide imputed wages globally, predictions are adjusted across countries according to Purchasing Power Parity as well as differences in wages across job categories between each country and the US. Finally, wages are adjusted for country-specific inflation. The imputed wages in our data behave similarly to observed wages (Amanzadeh et al., 2024).

A highlight of our data is that we observe granular location information for individual positions. When creating a position entry in a profile, users are asked to provide the location of that job. By observing this information, we are able to match individuals to specific cities around the world. An advantage of using the user-provided location information is that we are able to specifically

[4]Thus, while the wages are expected to capture information relevant to location, there is not a mechanical city effect added to the estimated wages.

identify where individuals are working, even if they work for a multi-establishment firm, including multinational firms. Our global coverage is extensive, with over 220,000 cities in our data. Figure 1b plots each city in our dataset with over 100 users by its average salary. While the large differences in income across countries make it difficult to identify within-country variation, Figures 2a and 2d highlight the richness of our city-level data. In addition to the large number of cities across both countries, these figures match the fact that we observe higher wages in wealthier regions of both countries, such as California and New England in the USA.

We take several steps to construct our primary sample of users for our analysis. For all work positions with no listed end date–which populate as "present" on LinkedIn–we assign an end date that is equal to the date Revelio Labs collected the data at the end of 2023. We then expand the data from the user-by-position level to user-by-year level by creating yearly entries based on the start and end date of positions. We apply a growth rate of 3%, which roughly matches the average year-on-year growth rate of salaries in the US, backward from the end date of the position for each previous year. Our results are robust to also simply using a constant wage and log interpolation between the estimated start and end salaries. Because our analysis only requires users that move between cities at some point in their careers, we then restrict our sample to all users that appear in two cities throughout our time period.[5]

3 Descriptive Facts

Figure 1: The Distribution of LinkedIn Users and Wages Across the World

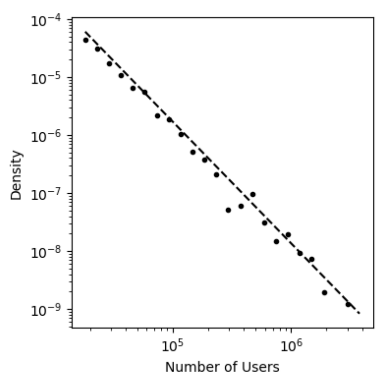

(a) Cities by Number of Users

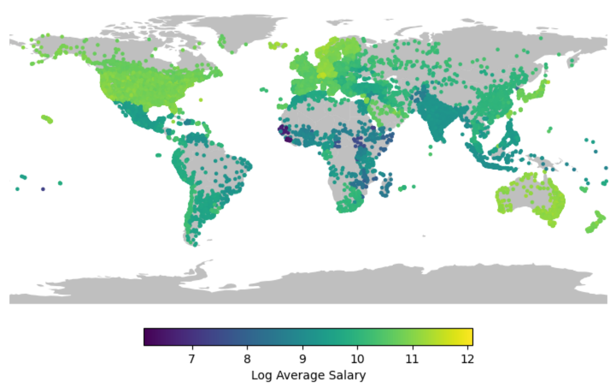

(b) Average Log City Wages, Global

Notes: The left panel shows the density of cities by the number of LinkedIn users. The x-axis is on a log scale and shows the number of LinkedIn users present in our database. The vertical axis is the probability that a city has a certain number of users. The right panel shows the average wages calculated in each city with over 100 users in our data (on a log scale), across the world.

[5]Our current sample excludes individuals who have worked in more than two cities during their careers. This is because we want to avoid concerns regarding return migration or falsely attributing wage gains to the wrong location based on the order of moves.

We describe certain key aspects of our data before estimating city effects. The distribution of users across cities follows recognizable patterns. Figure 1a shows that a large fraction of users are concentrated in a few major cities, while most cities have significantly fewer users, following a power-law distribution. This reflects the observed population distribution patterns across cities globally (Düben and Krause, 2021).

Figure 2: Geographic distribution of average city salaries by country

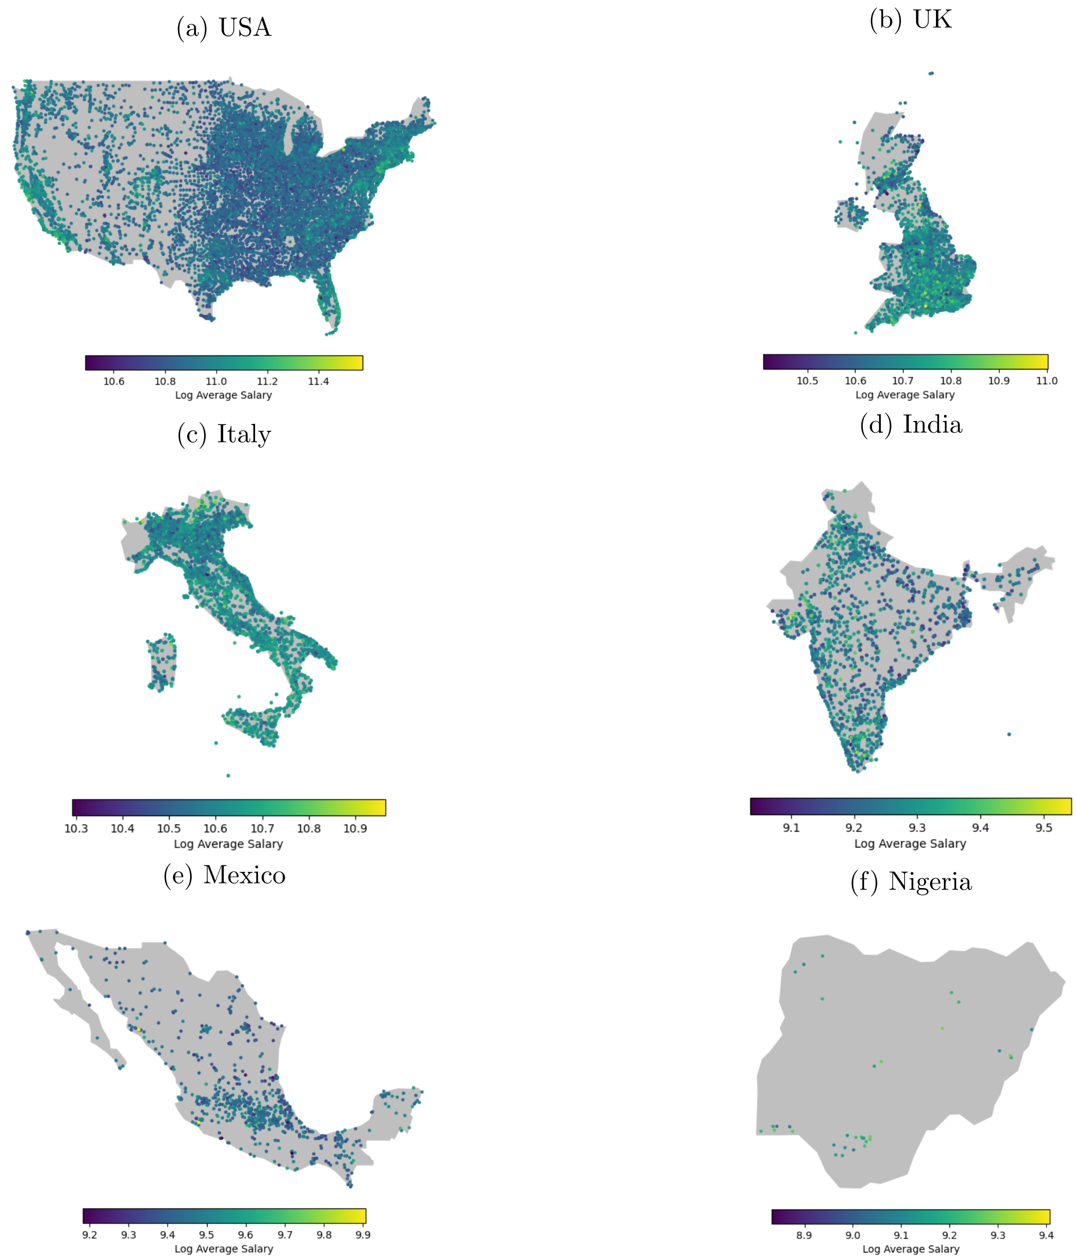

Notes: The maps plot average log(salaries) in each city. Legend scale varies by country. Sample restricted to cities with more than 100 LinkedIn users.

3.1 Average Wages by City

Figure 1b plots the distribution of average city-level wages in cities around the world. A few key patterns emerge. First, on the coverage, the database has wide coverage in all countries. The blank grey zones correspond with less populated parts of the globe. Second, there is large regional heterogeneity in wages, as one may expect, with the low end of the wage distribution appearing in Africa and South Asia, and the high end appearing in Europe. Third, country borders are starkly visible in certain parts, highlighting the possibly large gains from cross-border migration. This is particularly evident, for instance, on the border between Mexico and the US, or between North Africa and Southern Europe. And last, while a bit harder to see, there is meaningful variation within countries as well. For instance, even though almost all US cities have high average salaries, even in the US, there are many higher-earning cities.

To explore the within-country variation better, we examine certain countries in isolation, allowing the legend color range to vary by country. We show these maps in Figure 2, after picking countries across the development spectrum. The maps show the regional variance within countries, while also highlighting the density of our city coverage (despite restricting the sample to cities with more than 100 users). For instance, while there are relatively higher wages along the coasts than in the Midwest, there are pockets of cities in the Midwest with high-salaries as well. The UK seems to have more relatively high-wage cities in the south, where it is also relatively denser in terms of coverage. Italy's richest cities seem to agglomerate around Lombardy and the rest of the north, while the south has relatively lower wages. In India, the southern and western parts of the country have relatively more high-wage cities, with the central and eastern parts showing some lower-wage centers. Importantly, the range (on a log scale) varies greatly across the income spectrum.

3.2 Movers and Transition Matrices

Our estimation will rely on workers who move across cities. We describe this sample of movers, and the moves that they do. Figure 3a first shows the distribution of movers by origin income. There are more moves seen at higher ends of the city-income distribution. We divide cities into 5 quintiles to more easily examine transitions across location income quintiles.

In the next panel, Figure 3b shows the transition matrix of all recorded cross-city moves, with the origin city quintile being on the horizontal axis, and the destination quintile on the vertical axis. Most moves lie along the diagonal. This, perhaps, reflects the fact that most moves are within countries, and that country borders matter greatly for city-level income differences. Yet, there does seem to be some mobility from the fourth (origin) quintile to the fifth (destination) quintile as well, suggesting some potentially higher mobility at the upper middle-income range of origin cities.

Finally, in Figures 4a to 4c, we examine transitions within and across borders. First, in Figure 4a, we study within-country moves. The north-west triangle of the graph being brighter than the south-east suggests that most moves are to higher-wage cities, from lower-wage cities. That is,

Figure 3: Movers Sample: Income Quintiles and Transition Matrices

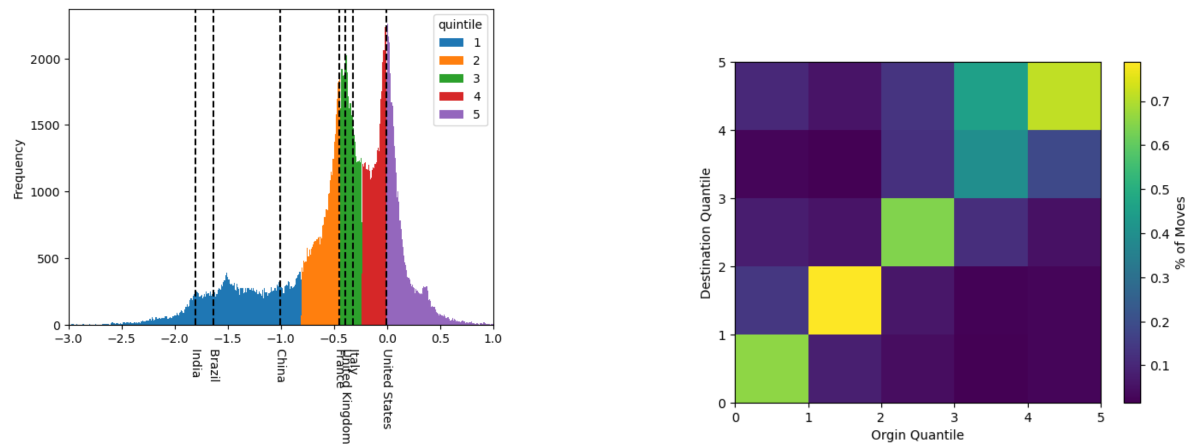

(a) Wage Distribution of Locations by Movers

(b) Transition Matrices (All Moves)

Notes: The left panel shows the density of movers by average city wages at the origin. We also show the distribution for 5 equally-sized city-level wage quintiles of the data. The right panel shows the joint distribution of moves by origin and destination city wage quintile.

individuals are less likely to move to a lower-wage city, within a country. The top right grid being distinctly higher density suggests, for instance, that those in the top quintile of origin cities are only more likely to move to other cities in the top quintile (and very unlikely to move to cities in the bottom quintile).

Figure 4b looks at the subset of moves that occur across country borders. The axes are the income quintiles within an origin or destination country. For instance, if a person migrates from India to the US, it plots the density for the income quintile of the Indian city, relative to all Indian cities on the origin x-axis, and the income quintile of the US city, relative to US cities on the destination y-axis. The high density in the top row suggests that when individuals cross borders, regardless of what type of city they come from, they choose to migrate to a high-income city in the destination country. That is, for international migrants, whether they come from low or high-income origins, their destination city is relatively wealthy compared to other cities in the destination country. In contrast, the bottom row being low density suggests immigrants are much less likely to choose relatively poorer cities among destination-country cities.

Finally, in Figure 4c, we look at the US transition matrix. This reflects the average within-country transition matrix, whereby individuals are more likely to transition up, and move to relatively higher-wage cities. Again, the high density on the top-right grid shows that those who leave quintile-5 cities are most likely to move to another quintile-5 city.

These patterns together suggest that individuals do seek out higher-wage cities. Below, we aim to better understand whether these cities have higher wages because high-ability people sort into them or because they make individuals more productive.

Figure 4: Transition Matrices for Migrants

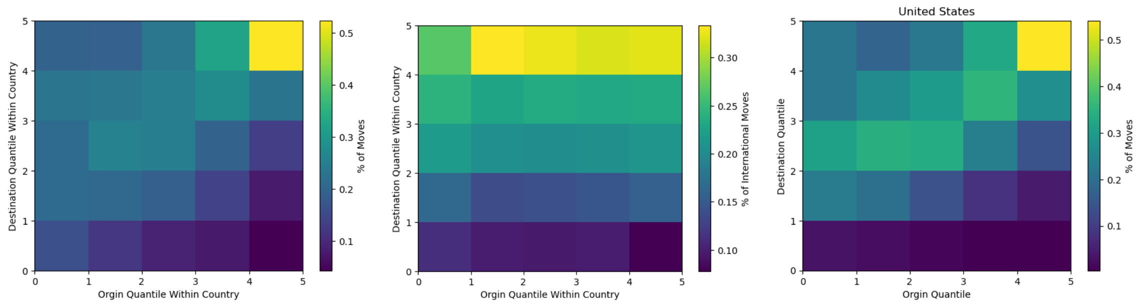

(a) Within Country Moves (b) Cross Border Moves (c) US: Transition Matrix

Notes: The figures show the transition matrices for within-country moves (left panel), cross-border moves (middle panel), and the US (right panel). The cross-border moves plot the income quintiles within a country. That is, for a person migrating from India to the US, it plots the density for the income quintile of the Indian city, relative to Indian cities (x-axis origin) and the income quintile of the US city, relative to US cities (y-axis destination).

3.3 Descriptive Event Studies

We first begin with studying bilateral moves across any two pairs of cities, in an event study framework. If wage variation across cities is primarily driven by sorting, individuals moving between cities would not necessarily experience systematic changes in their wages. Conversely, if city-specific wage premiums exist, individuals moving to cities with higher average wages for other workers would experience wage gains. In comparison, those moving to cities with lower average wages would face wage losses.

Specification: To investigate this empirically, we first define δ_i that denotes the difference in average log wages between the mover's destination and origin city. Specifically, for mover i, whose origin and destination cities are $o(i)$ and $d(i)$, respectively:

$$\delta_i = \bar{y}_{d(i)} - \bar{y}_{o(i)} \ , \tag{1}$$

where $\bar{y}_j$ denotes the average of $\bar{y}_{jt}$ across t, and $\bar{y}_{jt}$ is the expectation of outcome y_{it} across workers living in city j in year t. Following Finkelstein et al. (2016), the event-study specification that follows is:

$$y_{it} = \alpha_i + \tau_t + I_{r(i,t)} + \theta_{\mathrm{r}(i,t)}\delta_i I_{r(i,t)} + \eta_{it} \ , \tag{2}$$

where y_{it} is the log wage of individual i in calendar year t. α_i is an individual fixed effect that captures time-invariant skills of worker i and τ_t controls for calendar year fixed effects. $I_{r(i,t)}$ is a vector for relative-years, where for mover i who moves in year t^* relative year $r(i,t) = t - t^*$. The relative-year specific coefficients $\theta_{\mathrm{r}(i,t)}$ are our main parameters of interest. They capture changes in y_{it} around the move, scaled by δ_i.

Results: We present our event study results from Equation 2 for different populations of movers in our sample. On the y-axis, we plot estimated coefficients $\theta_{r(i,t)}$ for the corresponding relative year $r(i,t)$ on the x-axis. $\theta_{r(i,t)}$ captures changes in wages relative to the years surrounding the move, scaled by the difference in average wages between the mover's destination and origin cities. We normalize the value for $r(i,t)$ = -1 to 0.

Figure 5a plots estimated coefficients $\theta_{r(i,t)}$ for individuals whose origin and destination city are located in different countries ('international movers'). The figure shows a sharp discontinuous jump in the year of the move, from 0 to approximately 0.95. This implies a city-effect share of $\approx$ 95% in the observed variation in wages across cities internationally (Finkelstein et al., 2016), and suggests large potential wage gains for international moves, irrespective of individual-level skills.

Figure 5: Event Study for International and Internal Movers

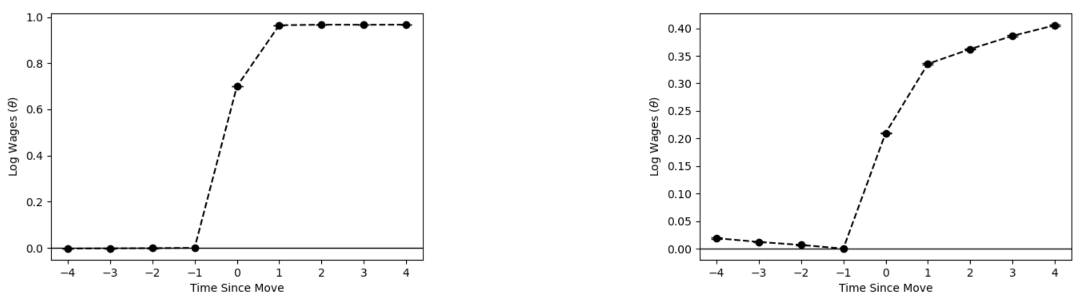

(a) International Movers, N = 2,736,283

(b) Internal Movers, N = 13,222,406

Both estimation samples are subset to only include individuals that worked in the year prior to moving ($t = -1$). Standard errors clustered at the user level.

Figure 5b shows the results for movers who move to cities within the same country as their origin. The jump from 0 to 0.35 around the move implies the city share of 35% in the observed variation in wages across cities in this sample. The sample for internal movers contains moves across 123,431 cities, while the international moves are across 56,422 cities.

While city effects are quite meaningful, the much larger effects for cross-border moves suggest country effects are extremely important, too. That is, moving from Bangalore to San Francisco is likely to have a much larger impact on one's earnings than moving from Omaha to San Francisco, partly because moving from India to the US also has a big impact on a worker's earnings.

Figure 6 and 7 plot the event study figures for a sample of developed and developing countries. The presence of some post-move trends implies that δ_i is positively correlated to wage growth after the move.

3.4 Average Gains by Origin-Destination Pair

The event study graphs show stark increases in earnings, as a function of the origin-destination wage differential. Keeping with the analysis of bilateral (origin-destination pair) moves, in this section,

Figure 6: Event Study for Internal Movers, Developed Economies

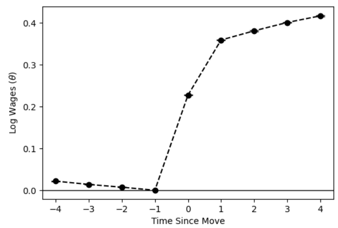

(a) USA, N = 5,869,926

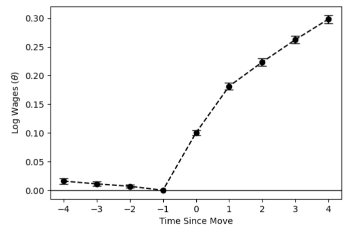

(b) UK, N = 832,197

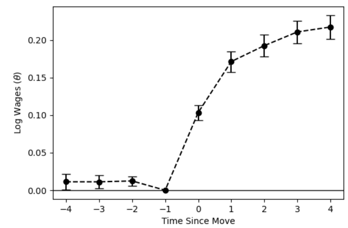

(c) Italy, N = 332,031

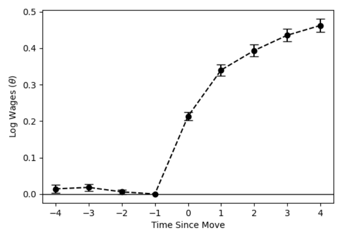

(d) Germany, N = 433,050

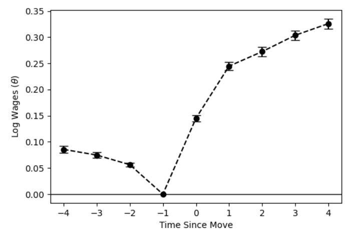

(e) France, N = 905,739

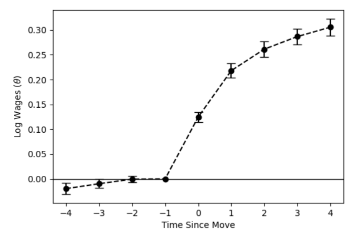

(f) Australia, N = 261,703

All estimation samples are subset to only include individuals that worked in the year prior to moving ($t = -1$). Standard errors clustered at the user level. Number of observations indicates number of users in each sample.

Figure 7: Event Study for Internal Movers, Developing Economies

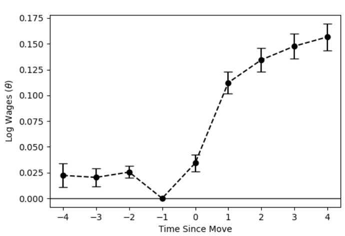

(a) India, N = 774,142

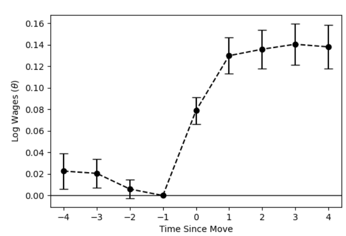

(b) Mexico, N = 202,888

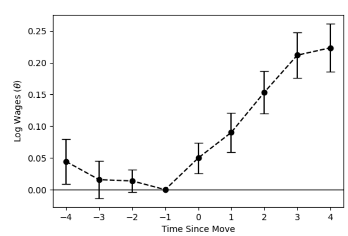

(c) Nigeria, N = 68,230

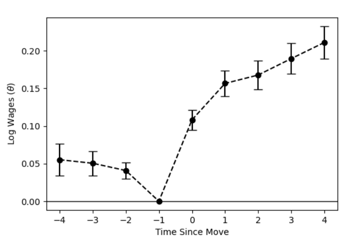

(d) Indonesia, N = 183,564

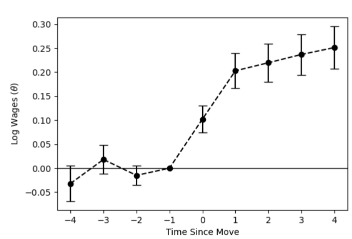

(e) Chile, N = 65,282

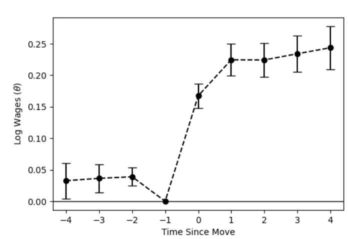

(f) Philippines, N = 148,944

All estimation samples are subset to only include individuals that worked in the year prior to moving ($t = -1$). Standard errors clustered at the user level. Number of observations indicates number of users in each sample.

we describe the relationship between the difference in wages in the two years prior to moving to wages in the two years after a move and the average wage differential at the origin-destination pair level ($\delta_i \equiv \bar{y}_{d(i)} - \bar{y}_{o(i)}$).

Figure 8: Change in Wages by Difference in Pairwise City Wages

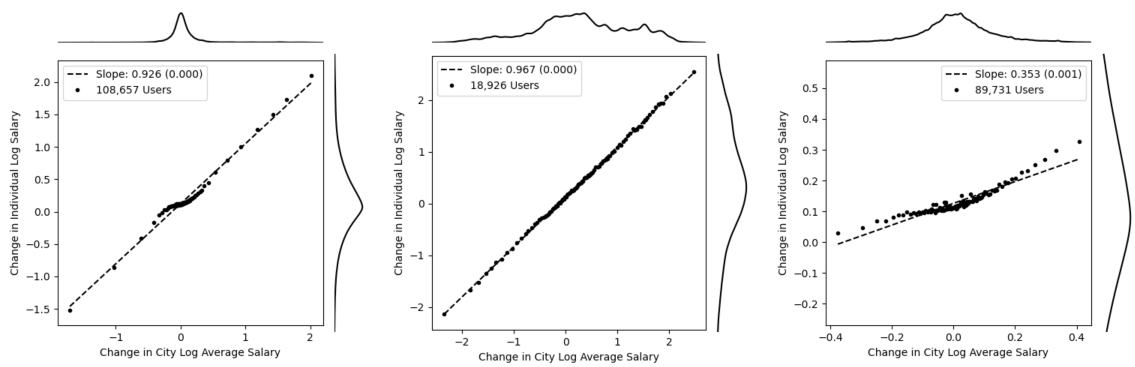

(a) Full Sample (b) International Moves (c) Internal Moves

We plot the difference in wages two periods before and after the move (y-axis) vs the difference in origin and destination average salaries corresponding to the move (x-axis). The left panel is for all moves. The middle panel for international moves. And the right panel for within-country moves. External axes plot the density of observations. Binned scatters with the number of unique users within each bin are reported.

This exercise is informative of how we go from the event-study estimates to the city effects. Essentially, the event study graphs plot an average jump in wages for individuals who move between cities. Yet, we may expect that this average hides a fair bit of heterogeneity, and non-linearities based on the origin-destination wage differentials. For instance, individuals moving between two rich cities may see a different proportionate gain (as a fraction of the wage differentials between the two cities), than individuals moving from a poor to rich city.

Figure 8a plots the relationship between the jump in wages for an origin-destination pair (y-axis), and the corresponding origin-destination average salary differential (x-axis). We plot the line of best-fit over this relationship, the slope of which should be similar to the size of the average jump in the event-study graphs. We also plot the density of observations on the axes.

The relationship in Figure 8a is roughly linear, with a flatter portion around the center. This may reflect the fact that most of the moves around the center of the graph are for within-country moves with smaller wage differentials, while the larger wage differentials reflect cross-border moves. Since country borders create stark wage differences (relative to within-country city differentials), moving across countries can lead to potentially larger gains. We investigate this relationship further in Figures 8b and 8c.

Indeed, the international moves (Figure 8b) show a stark (almost one-to-one) relationship with average city wage differences. The wage differential between Bangalore and San Francisco reflects the fact that, on average, when a worker moves between these cities, they will earn substantially more: almost the entire average wage difference. In contrast, for internal moves (Figure 8c), the relationship is a lot flatter (a slope of 0.34). Individuals that move between Omaha and San Francisco, will see a relatively smaller increase in their wages, as a fraction of the wage differential between the two cities.

Figure 9: Change in Wages by Difference in Pairwise City Wages

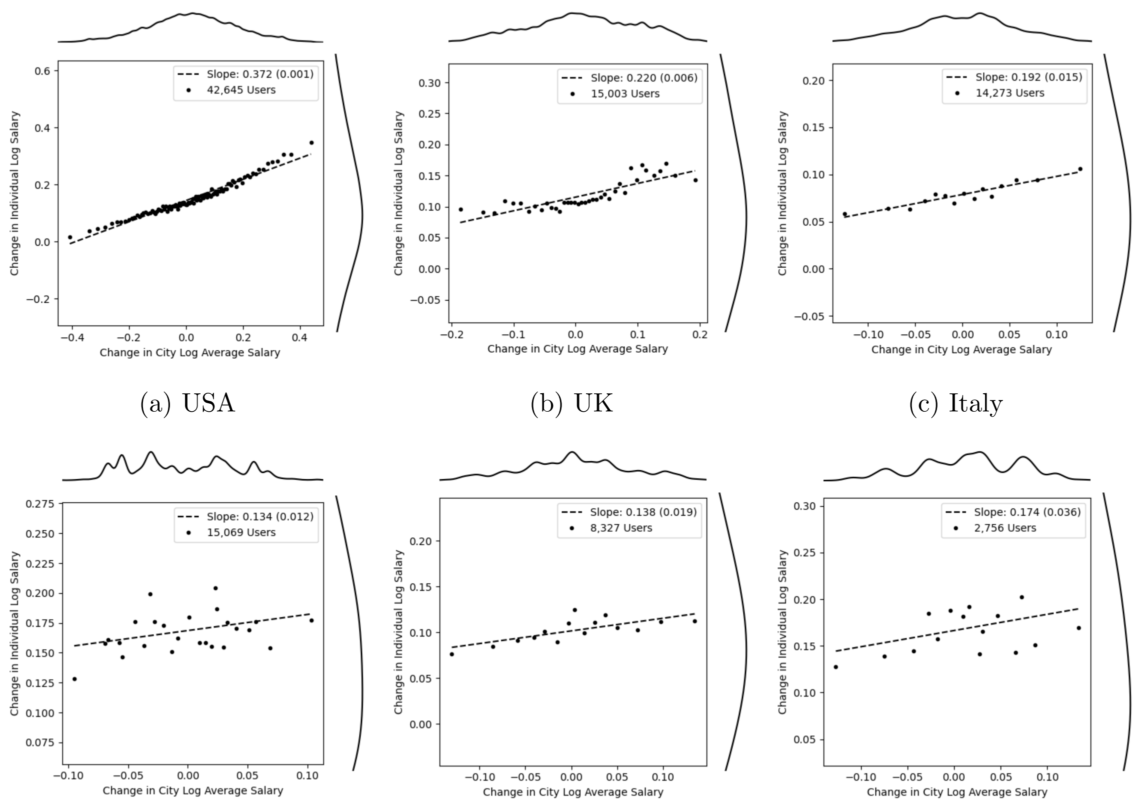

(a) USA (b) UK (c) Italy

(d) India (e) Mexico (f) Nigeria

We plot the difference in wages two periods before and after the move (y-axis) vs the difference in origin and destination average salaries corresponding to the move (x-axis). External axes plot the density of observations. Binned scatters with the number of unique users within each bin are reported.

These patterns reflect the fact that, given the (relatively) low migration barriers between Omaha and San Francisco, there has already been substantial internal migration whereby high-ability individuals have sorted to co-locate in one of the cities. That is, a fair amount of the wage differential between Omaha and San Francisco, reflects ability-based sorting. Furthermore, the steep relationship in Figure 8a is largely driven by cross-border moves.

Next, we conduct the same exercise, within countries. Figure 9a shows a slope of 0.37, reflecting a similar size of the event study jump on average. This slope (and so the event study coefficient) is much flatter for a country like India. Figure 9d shows a slope of about 0.13. This suggests, that within India, ability-based sorting may play a substantially more important role in driving pairwise city wage differentials. Someone who moves internally in the US, will see a substantially larger rise in wages (as a fraction of the pairwise city wage differential), than someone who moves within India. The differences in this relationship between Figures 9a-9c and Figures 9d-9f reflect the fact that the changes in salary associated with moving are typically higher in developed countries compared to developing countries.

Figure 10: Changes in Wages by Difference in City Wages, Over Time

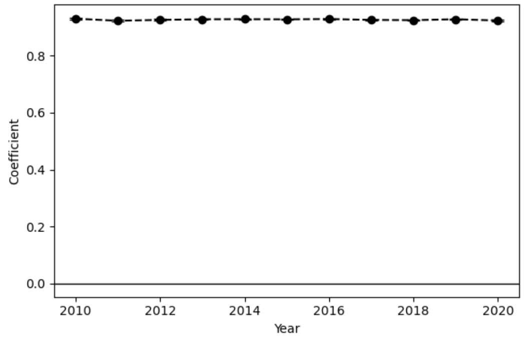

(a) International Moves

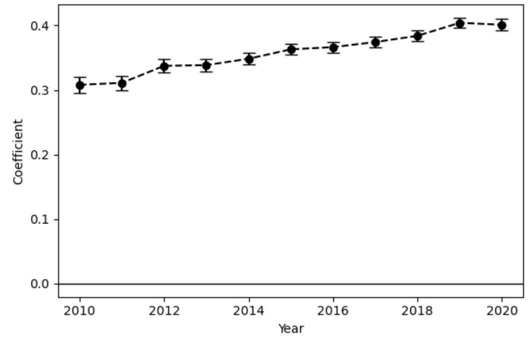

(b) Internal Moves

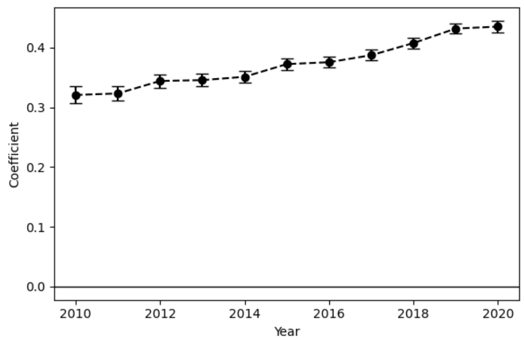

(c) USA Moves

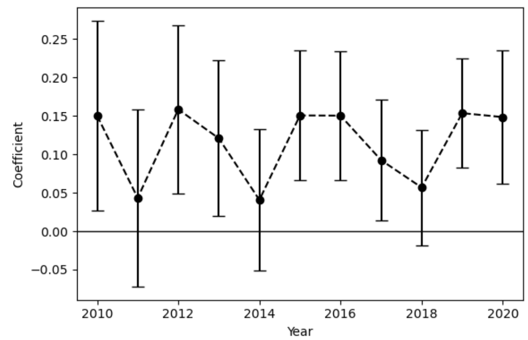

(d) India Moves

We plot the difference in wages two periods before and after the move (y-axis) vs the difference in origin and destination average salaries corresponding to the move (x-axis), by year of move. The left panel is for international moves, and the right panel is for within-country moves.

We also find that this relationship between wage gains and city wage differentials is not static over time. Figure 10 plots the same relationship between wages gains before and after a move against origin and destination wage differentials, but now disaggregated by the year of the move. While this relationship is constant for international moves, Figure 10b reveals that there has been a steady increase over time in the relationship between wage gains and city wage differentials for internal moves; between 2010 and 2020, we find that–for a move between locations with a given wage differential–the average increase in wages before and after a move has increased nearly 35%. This relationship could reflect the growing importance of city effects in determining productivity or increasing frictions, such as urban congestion and housing supply, that have limited the role that sorting by ability across space plays. Although we do see a clear trend when grouping all internal moves, it is not uniform across countries. Figures 10c and 10d plot this relationship for the United States and India, respectively. We find that the relationship for the US looks similar to the general trend; however, we find virtually no changes in the relationship between wage gains from a move and city differentials in India over time.

4 Empirical Strategy to Estimate City Effects

We now move away from bilateral pairwise differences in wages between origins and destinations, to isolate the overall city effects for each city. To set ideas, let us consider the standard specification in the mover's design literature (Abowd et al., 1999; Finkelstein et al., 2016). We first begin with a simple two-way fixed effects model such that log wage y_{it} of individual i in year t is the sum of a worker component α_i, a city component $\psi_{\boldsymbol{J}(i,t)}$, time-varying characteristics $(\tau_t, x'_{it}\beta)$, and an error component ϵ_{it}:

$$y_{it} = \alpha_i + \tau_t + \psi_{\boldsymbol{J}(i,t)} + x'_{it}\beta + \epsilon_{it}. \tag{3}$$

The function $\boldsymbol{J}(i,t)$ indicates the city where worker i was employed in year t. α_i captures time-invariant individual-specific characteristics, τ_t controls for calendar year fixed effects. $x'_{it}\beta$ controls for time-varying characteristics, including the year relative to move. The inclusion of relative year effects accounts for the possibility that the decision to move is correlated with wage shocks – for instance, when laid-off workers relocate to seek employment.

To fix ideas, it is useful to explore what this basic empirical setup allows for and does not allow for. Identification arises only from individuals moving cities, and only for the set of cities indirectly connected by individual movements (the largest connected set). The framework is flexible in that it accounts for certain types of systematic patterns of city choice. It allows for systematic mobility related to time-invariant individual characteristics or city characteristics. That is, individuals may be more likely to move from low-wage/low-amenity to high-wage/high-amenity cities (these city features can be time-varying as well). Furthermore, more productive workers may be more likely to switch cities, and our individual-level fixed effects account for that. Together, these features also allow for assortative matching (high productive workers move to high productive cities).

Yet, to identify effects, one needs to rely on two standard assumptions from the AKM literature. First, we assume that individual effects (α_i) and city effects (ψ_J) are additively separable in logs, precluding any interactions between them. This is an attractive assumption as it is consistent with a simple log-linear production function (so multiplicative in levels) of worker effects and city effects.

Second, we assume exogenous mobility, meaning that conditional on individual and city fixed effects (and other time-varying controls), the error term ϵ is uncorrelated with changes in $J(i,t)$. Three forms of mobility could violate this assumption (Card et al., 2013). The first is sorting based on city-individual match quality (Roy, 1951); if such sorting occurs, the interpretation of our city effects would change because different workers would receive different wage premia within the same city depending on match quality. While assortative matching does not affect the interpretation of our estimates, sorting on match quality would.

The second potential violation of the exogenous mobility assumption arises from 'drift' in individual fixed effects (e.g., due to employment shocks or human capital accumulation) that correlates with worker mobility across city years. People who lose their jobs or get demoted, may be more likely to move cities. Or, obtaining outside offers from other cities may bid up wages in one's current city, followed by a move. Similarly, it may take time to adapt to a city (which increases wages over time), and those who adapt better/faster may also be more (or less) likely to move. Alternatively, workers who struggle to adapt and see their wages decline may be more likely to leave a city. Similarly, if the move was driven by a one-time signing bonus, rather than a permanent wage increase, we would see a spike in the event study following the move, and a reduction thereafter. Relatedly, mobility correlated with transitory factors, such as seasonal variations or wage fluctuations, would also be a problem for our estimation.

Implementing tests proposed by Card et al. (2013); Finkelstein et al. (2016), we find that both additive separability and exogenous mobility are likely to hold in our setting.

First, we find evidence that endogenous mobility due to drift and transitory factors is unlikely to be an issue in our setting through our event-study analysis (e.g., Figures 6, 7 and 5). Particularly, we do not see any meaningful pre-trends to suggest drift in the portable component of the individual's earnings power. Neither do we see an 'Ashenfelter's dip' before a move (a negative job shock), or a one-time spike just after the move (a signing bonus).

Importantly, the event study framework allows for the fact that movers may be different from non-movers in levels of worker productivities (with individual fixed effects), and trends in wages around the moves (with individual time-since move fixed effects). Indeed, including these relative year fixed effects allows for the possibility that the decision to move is correlated with possible wage shocks or trajectories.

Furthermore, if mobility is endogenous due to idiosyncratic match quality, and the additive separability assumption does not hold, then the effect of moving between different types of cities (say ψ_A and ψ_B) should be asymmetric. That is, if worker i moved from City A to City B, the gain in wages should mirror the fall for workers who move from B to A.[6]

We test for this symmetry in Figures 8 and 9. We plot the average wage changes experienced around moves by individuals (y-axis) against the difference in average city wages (x-axis). The points on the left side of the x-axis (< 0) capture moves from a high-wage to a low-wage city, while the points on the right side (> 0) capture moves from low-wage to high-wage cities. We observe that wage changes associated with moves from low to high-wage locations are symmetric to those associated with moves from high to low-wage locations (Finkelstein et al., 2016). This provides strong support for the additive separability assumption and the absence of idiosyncratic match

[6]Specifically, if worker i moves from City A to City B:

$$\begin{aligned}&\mathbb{E}\left[y_{it} - y_{i(t-1)} \mid J(i,t) = B, J(i,t-1) = A\right] = \\ &\quad = \psi_B - \psi_A + \mathbb{E}\left[\varepsilon_{it} - \varepsilon_{i(t-1)} \mid J(i,t) = B, J(i,t-1) = A\right]\end{aligned}$$

while if the same worker moves in the opposite direction:

$$\begin{aligned}&\mathbb{E}\left[y_{it} - y_{i(t-1)} \mid J(i,t) = A, J(i,t-1) = B\right] = \\ &\quad = \psi_A - \psi_B + \mathbb{E}\left[\varepsilon_{it} - \varepsilon_{i(t-1)} \mid J(i,t) = A, J(i,t-1) = B\right] ,\end{aligned}$$

where the bias term is on the right-hand side. Without the bias term, the effect of moving should be symmetric ($\psi_B - \psi_A$ and $\psi_A - \psi_B$ here respectively).

components.

Taken together, the absence of pretrends and the symmetry in wage changes suggest that our model is well suited for estimating city effects on log wages.

4.1 Hierarchy Effects

The basic AKM framework, however, may provide estimates that are challenging to interpret when there is heterogeneity in the premiums paid by different firms within a city. For instance, if workers who move from low-wage cities come from above-average paying firms in their origin cities but migrate to high-wage cities and work in below-average paying firms, our estimated city effects would include what Card et al. (2023) refers to as a 'hierarchy effect.'

Depending on the study's objective, it may be worth separating out this effect. If the objective were to measure place effects that tell us how movers would do (acknowledging that immigrants to a particular city come from high-wage firms, but are more likely to work in lower-wage firms), then we may want part of the hierarchy effect. Alternatively, we may want to simply isolate the location effect in a hypothetical situation where we move individuals from random firms at an origin, and allocate them to random firms at a destination. Here, the 'hierarchy effect' would bias our estimates.

Fortunately, Card et al. (2023) provide an intuitive and tractable solution, that we implement. Instead of city-effects, we now estimate firm (establishment-level) effects, and aggregate the firm effects to measure the city effects. The specification we estimate is now:

$$y_{it} = \alpha_i + \tau_t + \gamma_{\boldsymbol{f}(i,t)} + x'_{it}\beta + \epsilon_{it} \ , \tag{4}$$

where $\boldsymbol{f}(i,t)$ is a function that indicates which establishment individual i was employed in, in year t. We define the city-specific wage premium to be a weighted average of the establishment effects within that city:

$$\Gamma_j = \frac{\sum_{j(f)=j} N_f \gamma_f}{\sum_{j(f)=j} N_f} \tag{5}$$

where $j(f)$ is a function giving the city for establishment f, γ_f is the establishment effect, and N_f is the number of person-year observations in our estimation sample for that establishment. The interpretation of this premium is that if a worker was randomly chosen and moved to a random firm at a destination, their earnings would increase by $\Gamma_j - \Gamma_{j'}$.

Importantly, the need to do this arises primarily because we may expect that the hierarchy effect is correlated with the city effect. That is, the conventional AKM provides meaningful estimates even if immigrants come from high-wage firms, and work in low-wage firms. However, if this difference in wage premia is systematically correlated with the productivity of the city, it would bias the estimates of a quasi-random thought experiment.

In practice, we limit our data to establishments with a FactSet Entity ID where at least five workers (in our dataset) work. In addition to the previous evidence supporting this design, the standard firm-level AKM assumptions are required for Equation 4, which have been stringently tested in various applications around the world (e.g., Card et al. (2013)).

We use this strategy as our main empirical strategy to estimate the city effects below.

5 City Effects

Figure 11: City Effects, Global

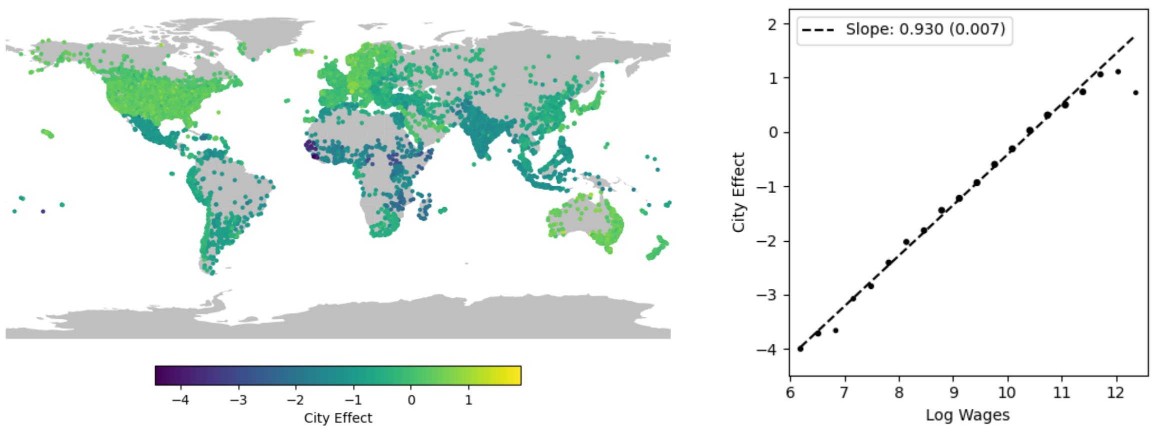

(a) Estimated City Effect

(b) Log City Wages vs City Effect

Panel (a) plots the city effects ψ_J estimated using Equation 3, using the same sample as the event study analysis in Figure 5a. Panel (b) plots the relationship between the Log City Effect and the Log Average City Salary, where points are equally spaced across the distribution and the size of each point is proportional to the fraction of cities in each wage bin. $N = 8,025,420$

We use our movers' sample and Equation 4 and 5 to estimate the city effects for each city around the world. To comprehensively understand the economic implications of these city effects, we conduct several analyses. First, we plot the estimated city effects for all cities worldwide based on international movers (Figure 11) and compare these with city effects derived from within-country movers (Figures 12, 13, 14). This comparison highlights the extent to which national borders influence city effects. To further disaggregate within-country differences, we examine regional variations in city effects, using the United States as a case study (Figure 15). Finally, we investigate the differences in the distribution of city effects across countries due to their implications for heterogeneity in potential returns to higher internal migration.

Collectively, these comparisons help us shed light on the importance of city effects under varying migration costs: international moves across countries, moves across cities in the same country (both within and across regions), and within-country moves for different countries. Subsequently, in Section 6, we explore what underlying economic primitives are reflected in our estimated city effects and the variance in city effects across countries.

Figure 12: City Effects, USA

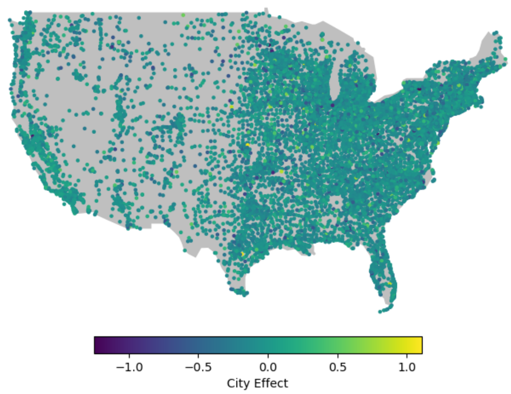

(a) Estimated City Effect

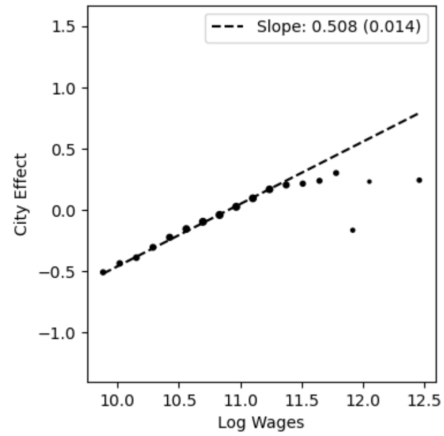

(b) Log City Wages vs City Effect

Panel (a) plots the city effects ψ_J estimated using Equation 3, using the same sample as the event study analysis in Figure 6a. Panel (b) plots the relationship between the City Effect and the Log Average City Salary weighting cities by the number of users, where points are equally spaced across the distribution and the size of each point is proportional to the fraction of users in cities in each wage bin.

5.1 City Effects Across and Within Countries

In Figure 11a, we plot our estimated city effects on a map. They do roughly reflect the map of average wages by city in Figure 1b, once again highlighting the importance of regions and country borders. Figure 11b plots the correlation between city effects and average wage premia. These are strongly positively correlated for cities across the world. A correlation of 0.93 indicates that cities play a significant role in explaining wage differentials across the globe, highlighting the importance of city-specific factors in shaping economic outcomes.

In Figures 12 and 13, we look at the example of only within-country internal moves for the US and India. The flatter relationship between city effects (y-axis) and average city salaries (x-axis) for the US ($\beta = 0.51$) and India ($\beta = 0.73$) suggests that city effects explain a smaller share of the city wage differentials within a country than international wage differentials ($\beta = 0.93$).

These results are consistent with our event study results from Figure 5, where potential wage gains from international migration dwarf the potential gains from internal migration. Additionally, the varying slopes observed in the country-level graphs suggest that the extent to which city effects explain wage differentials is not uniform across countries. These differences imply that the potential economic benefits of internal migration, particularly in terms of wage gains, depend significantly on the country-specific context. For instance, in countries where city effects play a larger role in determining wages, internal migration may offer greater opportunities for wage gains. We investigate the distribution in city effects across countries systematically in subsequent sections.

Figure 13: City Effects, India

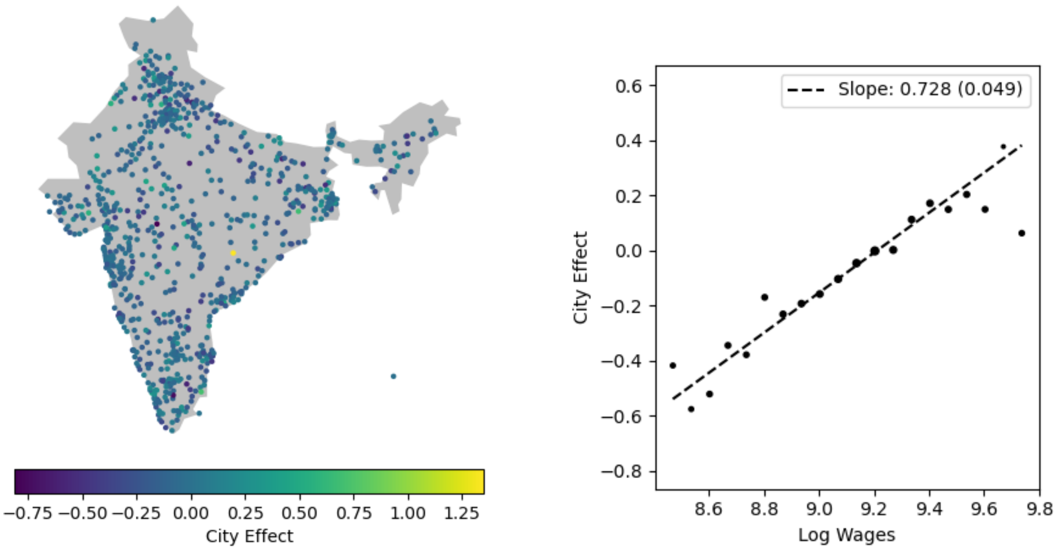

(a) Estimated City Effect

(b) Log City Wages vs City Effect

Panel (a) plots the city effects ψ_J estimated using Equation 3, using the same sample as the event study analysis in Figure 7a. Panel (b) plots the relationship between the City Effect and the Log Average City Salary weighting cities by the number of users, where points are equally spaced across the distribution and the size of each point is proportional to the fraction of users in cities in each wage bin.

5.2 The Contribution of City Effects to Wages

We report additive decomposition of the difference between high and low-wage areas, focusing on the share of the wage difference between two regions that are explained by city effects (Finkelstein et al., 2016). We define the share of wage difference between area R and R' explained by city effects as:

$$S_{city}\left(R, R'\right) = \frac{\Gamma_R - \Gamma_{R'}}{\bar{y}_R - \bar{y}_{R'}}\ , \tag{6}$$

where Γ_A is the average city effect for cities in area A, and $\bar{y}_A$ is the average log wage for cities in area A. Table 1 reports the decomposition of the variation in city wages across the world. Table 2 and Table 3 report the decomposition for a sample of developed countries (United States, United Kingdom, and Italy) and developing countries (India, Mexico, and Nigeria), respectively.

Column (1) breaks down the difference between cities with above-median and below-median salaries. Column (2) shows the difference between cities in the top and bottom quartiles, and Column (3) shows the difference between top and bottom 5%. Columns (4) looks at the differences between specific cities.

Additionally, in the spirit of Card et al. (2023), we can average Equation 4 across workers and years, to decompose differences in mean wages across cities:

$$\begin{aligned} \bar{y}_j &= \bar{\alpha}_j + \Gamma_j + \bar{x}_j\beta, \\ \mathrm{Var}(\bar{y}_j) &= \mathrm{Var}(\bar{\alpha}_j + \Gamma_j + \bar{x}_j\beta) = \mathrm{Cov}(\bar{y}_j, \bar{\alpha}_j + \Gamma_j + \bar{x}_j\beta), \\ &= \mathrm{Cov}(\bar{y}_j, \bar{\alpha}_j) + \mathrm{Cov}(\bar{y}_j, \Gamma_j) + \mathrm{Cov}(\bar{y}_j, \bar{x}_j\beta). \end{aligned} \tag{7}$$

As such, the share of variance in mean city wages attributable to city effects is:

$$\Omega \equiv \frac{Cov(\bar{y_j}, \Gamma_j)}{Var(\bar{y_j})} \tag{8}$$

Ω is also equivalent to the slopes reported in Figures 11, 12, and 13. We report Ω in the last row of Column (5), estimated with the help of a simple regression of city effects on average wages ($\Gamma_j = a + \Omega\bar{y}_j + e_j$).

In Table 1, we find that 93% of the global variation in wages across cities is explained by city effects, and 87% of the difference in average wages among high-wage and low-median places is attributable to city effects. In reference to our example in the introduction, we find that 93% of the difference in wages between Bangalore and San Francisco is due to the difference in city effects.

For our sample of developed countries, we find that 45-51% of the variation in wages across cities can be attributed to city effects (Table 2). 45-57% of the differences in wages between below and above-median cities in our sample of developed countries are due to city effects. For the US, we find that city effects explain about 51% of the variation in wages across cities, similar to 50% found by Card et al. (2023) for variation across commuting zones. Additionally, regarding the difference between below and above-median-wage cities, differences in city effects explain 45% of the variation. For our sample of developing countries, we find that 48-73% of the difference in wages is explained by city effects (Table 3). Overall, city effects explain a meaningfully large share of variation in wages on average, and across different partitions. Additionally, the share of variation in wages explained by city effects tends to be higher in developing countries - something we test formally in Section 6.

5.3 Regions Within Countries

Regions play a crucial role in shaping wage differentials. This is evident, for example, when comparing the coastal United States to the Midwest, where each region exhibits distinct income levels. Migration patterns also tend to be concentrated within regions, leading to two important implications. First, sorting of individuals is likely to occur within regions, rather than across regions, which leads to heterogeneous relationships between city effects and average salaries. Second, this contributes to a distinctive relationship between city effects and average salaries: within regions, this relationship is relatively more flat, but when observed across the entire country, it becomes significantly steeper.

Table 1: Decomposition of Wage Differences: Global

	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Bangalore to San Francisco (4)	Var $(\bar{y}_j)$ (5)
Difference in Ln Wages	0.75	1.27	2.39	1.1	–
Difference due to City	0.65	1.07	2.05	1.02	–
Share due to City	0.87	0.84	0.86	0.93	0.93

Notes: This table presents the decomposition of wage differences across cities globally. The first row computes the difference in average wages between the destination and origin. The second row reports the difference in average city effects between the destination and origin region. The third row reports the share of difference in average wages between two areas due to city effects, which is the ratio of the second row to the first row (Equation 6).

Figure 14: City Effects, Other Countries

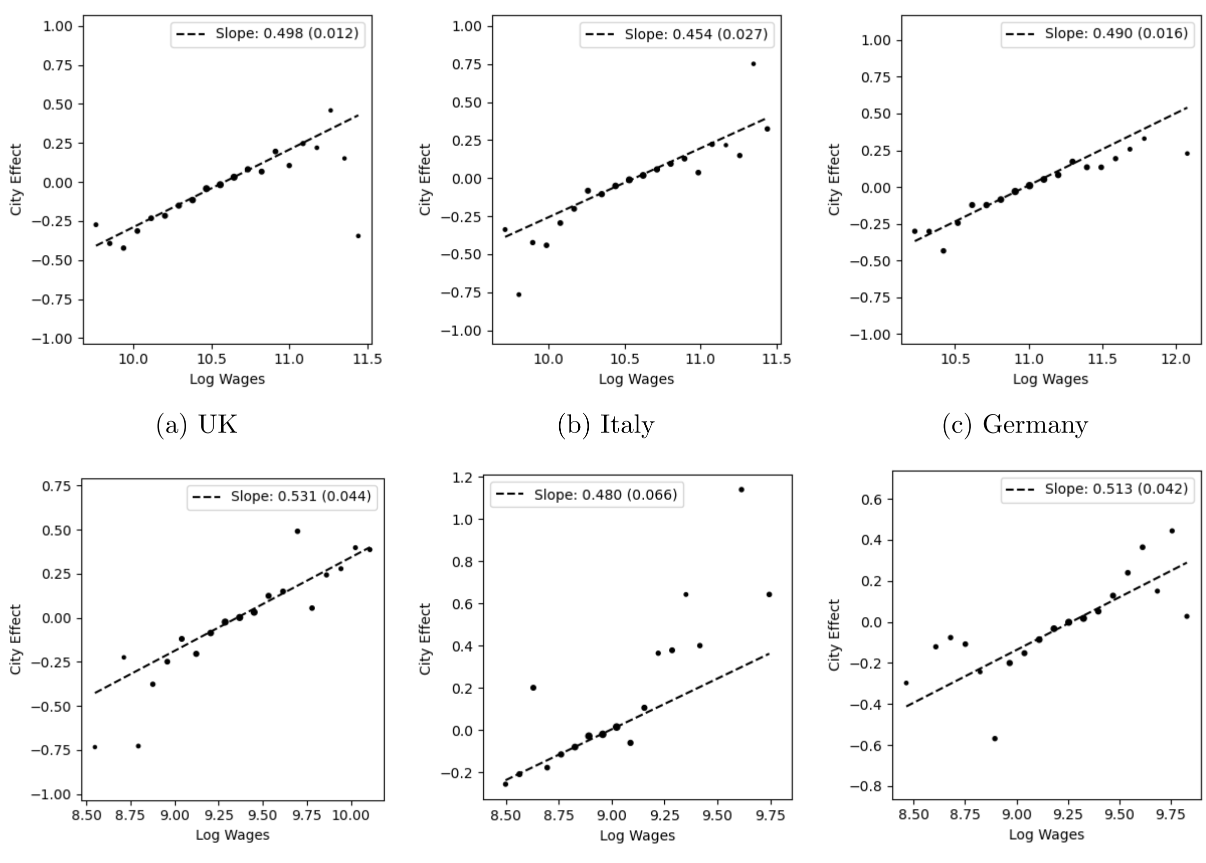

(a) UK (b) Italy (c) Germany

(d) Mexico (e) Nigeria (f) Philippines

Each panel plots the relationship between the City Effect and the Log Average City Salary weighting cities by the number of users, where points are equally spaced across the distribution and the size of each point is proportional to the fraction of users in cities in each wage bin.

Table 2: Decomposition of Wage Differences: Developed Countries

Panel A: United States	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	San Diego to New York (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.32	0.52	1.0	0.2	–
Difference due to City	0.14	0.24	0.45	0.04	–
Share due to City	0.45	0.45	0.45	0.22	0.51

Panel B: United Kingdom	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Brighton to London (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.23	0.4	0.79	0.12	–
Difference due to City	0.13	0.23	0.4	0.05	–
Share due to City	0.57	0.58	0.51	0.4	0.5

Panel C: Italy	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Florence to Rome (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.24	0.42	0.83	0.03	–
Difference due to City	0.11	0.2	0.42	0.02	–
Share due to City	0.45	0.47	0.5	0.65	0.45

Notes: This table presents the decomposition of wage differences across cities in different developed countries. In each panel, the first row computes the difference in average wages between the destination and origin region. The second row reports the difference in average city effects between the destination and origin region. The third row reports the share of difference in average wages between two areas due to city effects, which is the ratio of the second row to the first row (Equation 6).

Figures 15a and 15b illustrate this pattern. Regions that are geographically smaller and more integrated–such as the Northeast and West–have a closer relationship between city effects and average salaries. Additionally, there are significant level differences between the city effects across regions; the wealthiest, more urbanized regions of the US also have higher average city effects. Also, within each region of the US except one, the relationship between city effects and average salaries is more flat than the overall relationship in the US. However, because some regions are systematically wealthier than others, the relationship for the entire country becomes steeper, as reflected by the dotted black line.

Table 3: Decomposition of Wage Differences: Developing Countries

Panel A: India	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Kolkata to Bangalore (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.26	0.44	0.84	0.13	–
Difference due to City	0.17	0.29	0.47	0.1	–
Share due to City	0.65	0.65	0.56	0.74	0.73

Panel B: Mexico	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Tulum to Mexico City (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.24	0.41	0.79	0.12	–
Difference due to City	0.15	0.26	0.56	0.15	–
Share due to City	0.61	0.63	0.71	1.22	0.53

Panel C: Nigeria	Below to Above Median (1)	Bottom to Top 25% (2)	Bottom to Top 5% (3)	Benin City to Lagos (4)	Var ($\bar{y}_j$) (5)
Difference in Ln Wages	0.23	0.39	0.84	0.11	–
Difference due to City	0.13	0.2	0.68	0.03	–
Share due to City	0.55	0.5	0.8	0.25	0.48

Notes: This table presents the decomposition of wage differences across cities in different developing countries. In each panel, the first row computes the difference in average wages between the destination and origin region. The second row reports the difference in average city effects between the destination and origin region. The third row reports the share of difference in average wages between two areas due to city effects, which is the ratio of the second row to the first row (Equation 6).

These aggregation phenomena are likely present in other parts of the world as well, particularly in countries where some regions are consistently richer than others, and migration occurs predominantly within regions. This observation also helps explain why the slope of the relationship is steeper in our global analysis compared to the within-country analysis.

5.4 Validating the AKM specification

As discussed earlier in Section 4, in order for our city effect estimates to be unbiased, we need exogenous mobility to hold: moves across cities or firms have to be uncorrelated with the error

Figure 15: City effects by region, USA

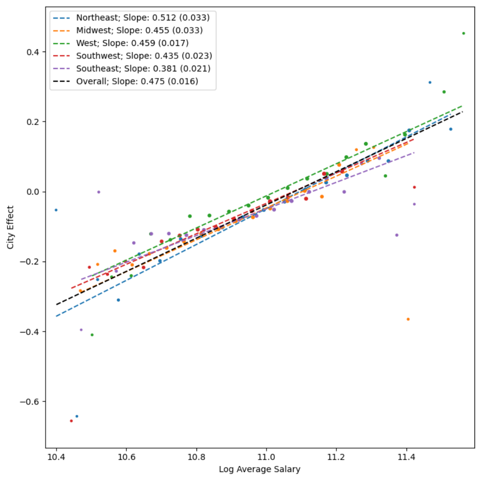

(a) Effects by US Regions

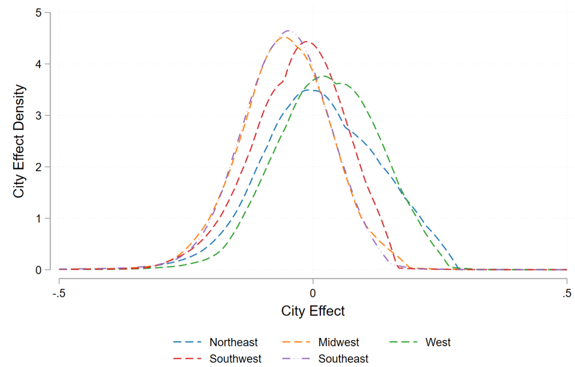

(b) Distribution of City Effects by US Regions

Panel (a) plots the relationship between city effects and log average salary by region of the US, and the overall fit line. Panel (b) plots the distribution of city effects, weighted by the number of users, by region of the US. Cities with estimated city effects $> \pm.5$ are excluded from the figure.

term ϵ_{it}. Our event-study analysis provides evidence to support exogenous mobility. Here, we conduct additional tests using our estimated city effects, as suggested in Card et al. (2023).

Panel A of Figure 16 plots the changes in time-adjusted earnings for movers across cities against the origin-to-destination average change in city effects. The change in adjusted earnings is calculated as the difference between adjusted earnings two years after and two years before the move. The 45' line corresponds to changes in earnings that would match the change in city effects. The scatter plot has a slope of 0.786, suggesting that on average, movers see 78.6% of the change in earnings that we would predict based on the changes in city effects when they move.

We can decompose the change in earnings for a worker who moves across period t and $t-1$ as follows:

$$y_{it} - y_{i,t-1} = \left(\Gamma_{c(i,t))} - \Gamma_{c(i,t-1)}\right) + \left(h_{f(i,t)} - h_{f(i,t-1)}\right) + \left(\epsilon_{it} - \epsilon_{it-1}\right) \ , \tag{9}$$

where Γ_c are the city effects, h_f is the hierarchy component of earnings (based on the firm's position in the local job ladder), and ϵ_{it} is the AKM residual.

If changes in the AKM residual are systematically related to the changes in city effects, it would suggest violations of the exogenous mobility assumption. We test this in Panel B of Figure 16. We see a relatively flat plot, with a precisely estimated slope of just 0.063. The small magnitude of the coefficient implies that any departure from the assumption is minimal.

We also plot the changes in the hierarchy component against the change in city effects in Panel C of Figure 16. The negative correlation indicates that workers who move from a low-city effect origin to a high-city effect destination, experience a decrease in the hierarchy effect. Thus, the firms they move to are lower on the local job ladder than the firms they came from. While workers who move in the opposite direction land at firms that are higher on the local job ladder.

Lastly, in Panel D we find that hierarchy adjusted earning changes are perfectly correlated with city effects (slope 1.06). Thus, similar to Card et al. (2023), we conclude that violations of exogenous mobility are quantitatively small, and our AKM specification captures city effects accurately. We make similar conclusions for the set of all movers (Figure 17) and city effects estimated using the connected set of international movers, except that the hierarchy effects do not matter quantitatively when comparing cities across countries.

Figure 16: Change in earnings components of city movers around move, against the change in city effect (within country movers)

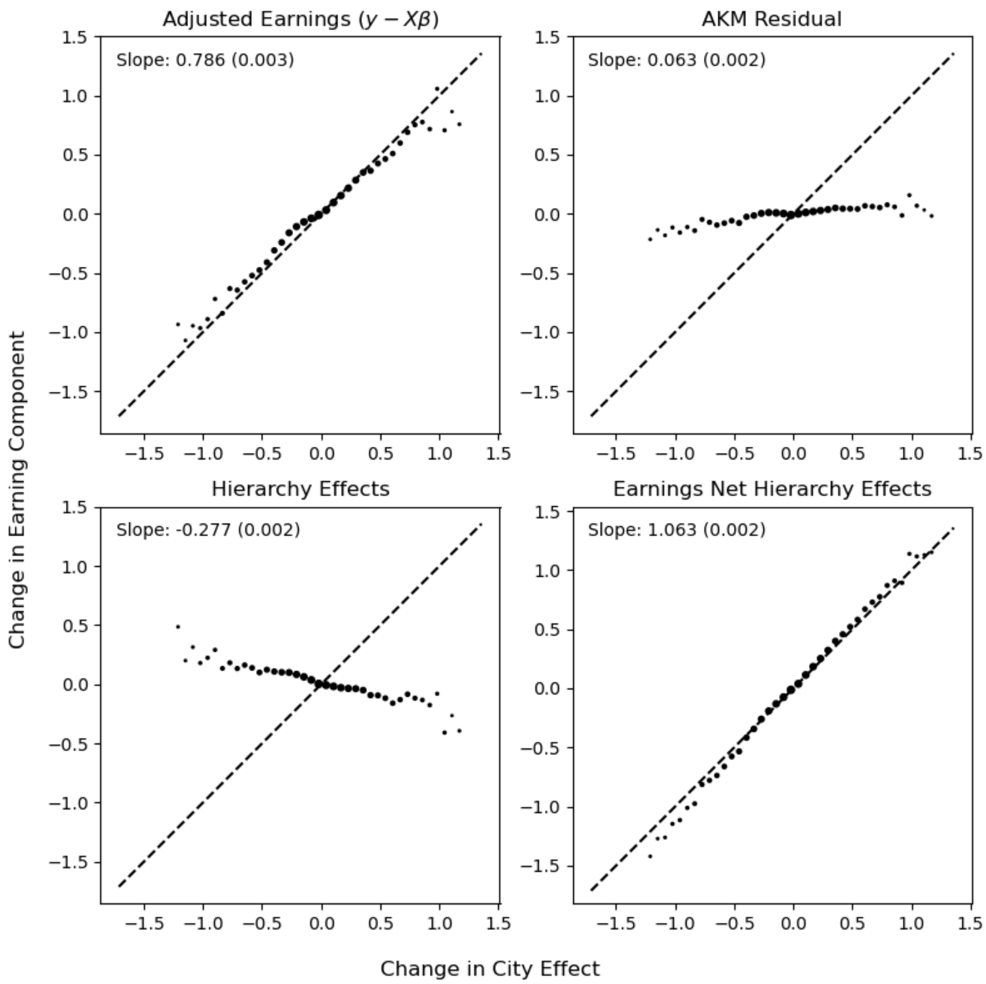

Notes: We look at changes in earnings two years before and after a move. X-axis plots the change in city effects for movers as destination city effect - origin city effect. City effects are constructed using equation 4. The slope is the best linear fit line. The sample is restricted to movers who move once and within their countries.

Figure 17: Change in earnings components of city movers around move, against the change in city effect (all movers)

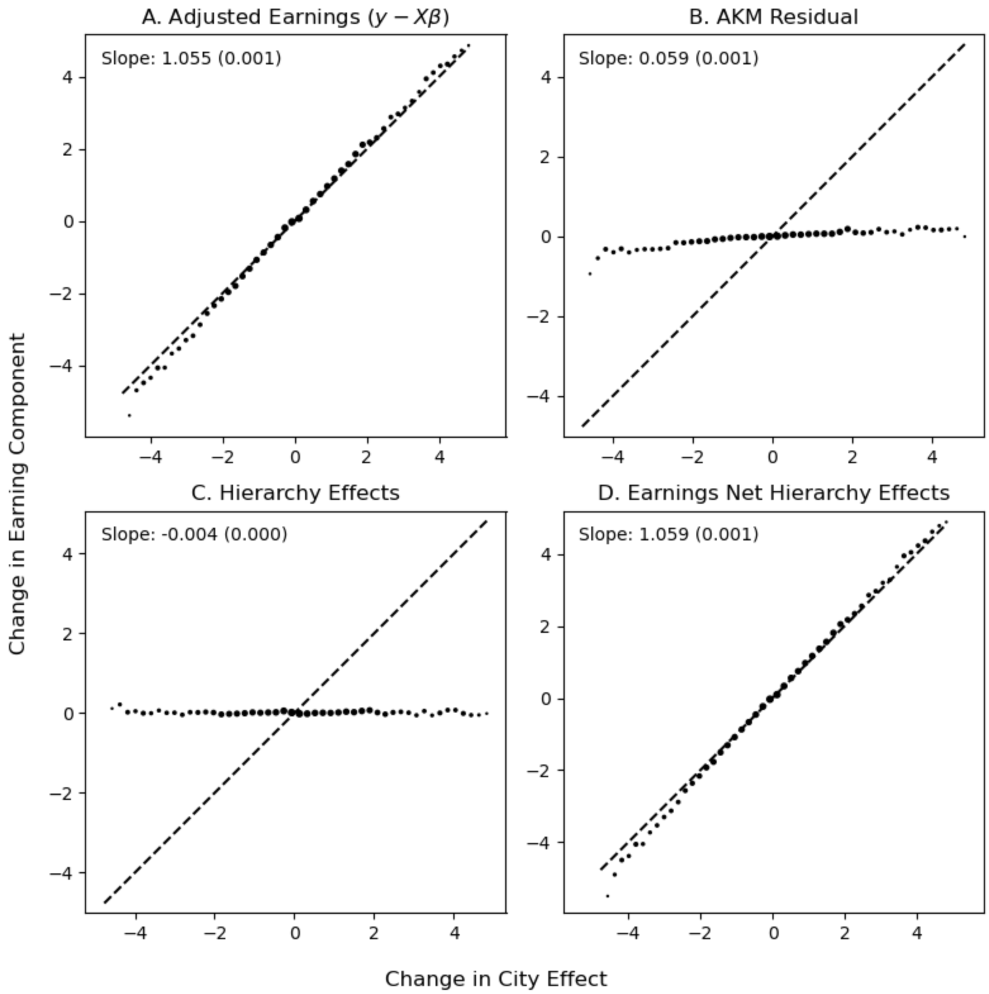

Notes: We look at changes in earnings two years before and after a move. X-axis plots the change in city effects for movers as destination city effect - origin city effect. City effects are constructed using equation 4. The slope is the best linear fit line. The sample is restricted to movers who move once.

6 What Explains the City Effects and Their Distribution

6.1 Explaining the City Effects

Figure 18a plots the standardized coefficients from our estimation while Figures 18b- 18f present binned scatter plots for each relationship. Consistent with conventional theories of agglomeration economies, we find that cities that are larger and have more modern firms (proxied by remote suitability) have higher estimated city effects. These patterns reflect other work that argues how city density and large populations facilitate knowledge sharing, with spillovers that drive city incomes (Duranton and Puga, 2004; Rosenthal and Strange, 2004).

Next, in Figure 18b, we find that cities with a lower industrial HHI have a wider variety of

Figure 18: Correlates with City Effects

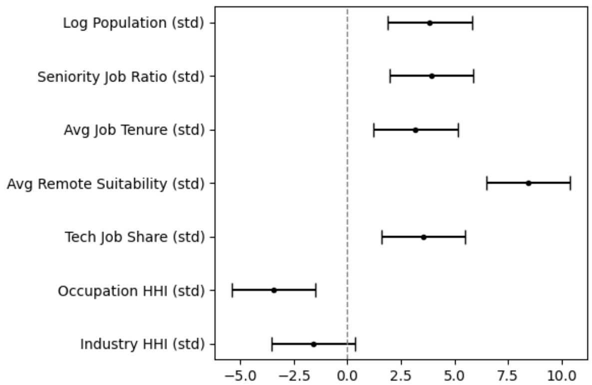

(a) Correlates of City Effects

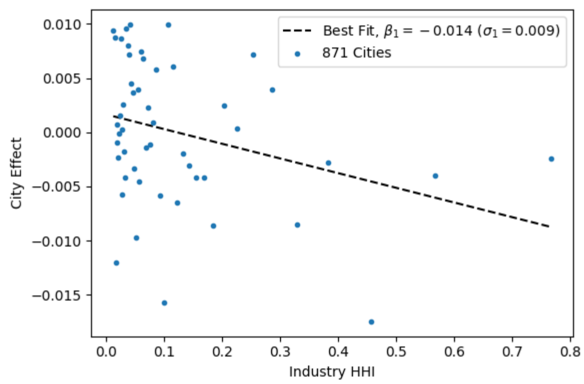

(b) Industrial Diversity

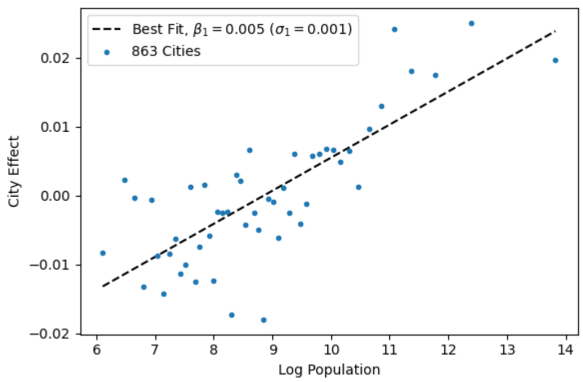

(c) Population

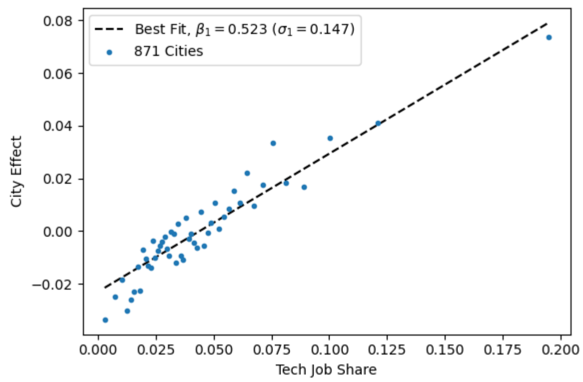

(d) Tech Job Share

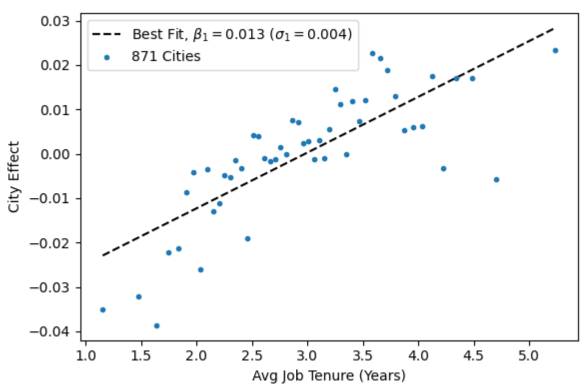

(e) Average Job Tenure

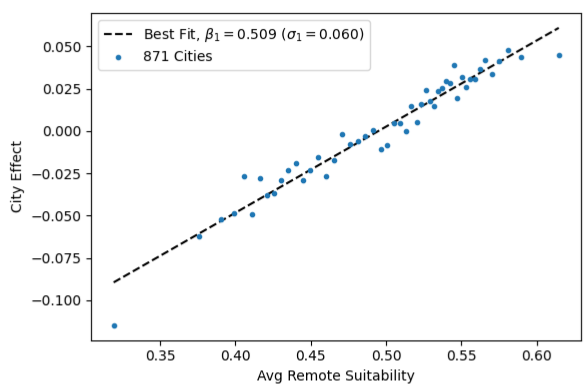

(f) Job Remote Suitability

Notes: We plot the city-level correlation between city characteristics and city effects with the inclusion of region fixed effects. Industrial diversity is the Hirschmann Herfindahl Index (HHI) of industry-wise employment. Log(population) is the city's population. Tech Job Share is the ratio of tech job titles as classified by ONET codes over total jobs in a city. Average job tenure is the average number of years workers stay at a job. Job remote suitability is a measure of whether the job-title-by-industry is conducive to remote work (based on job postings of remote work). Standard errors are clustered at the country level.

Figure 19: Correlates with City Effects By Country Characteristics

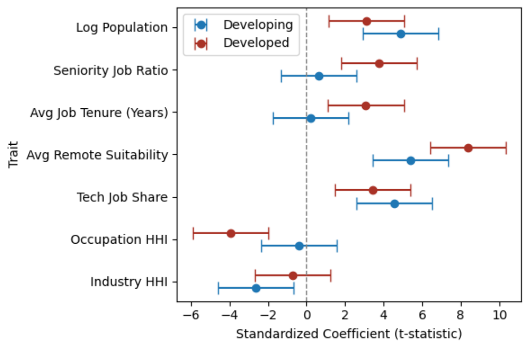

(a) By Development Status

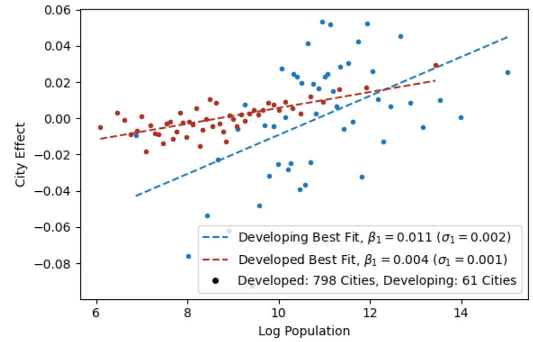

(b) Population

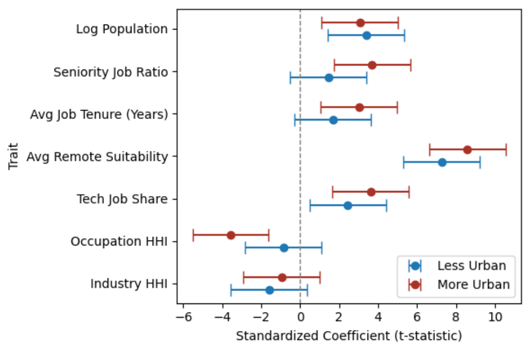

(c) By Urbanization

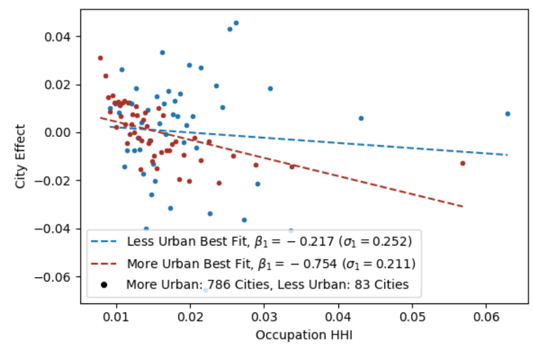

(d) Occupation HHI

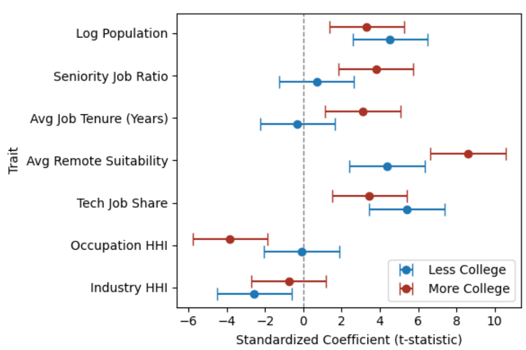

(e) By College Attainment

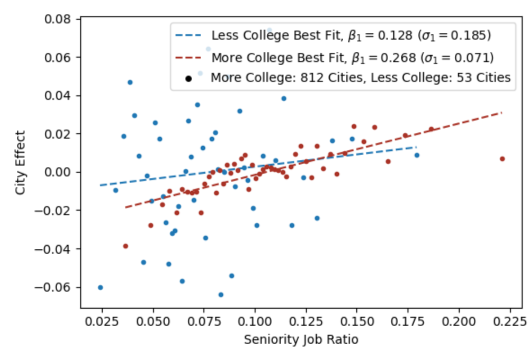

(f) Seniority Ratio

Notes: We first split our sample between countries with above/below median values for (i) GDP per capita (ii) urbanization rates and (iii) college attainment rates. We then plot the city-level correlation between city characteristics and city effects with the inclusion of country fixed effects for each sample split in the first column. The second column displays the relationship between population and city effects by sample split, where each point represents an equally sized number of cities. Industrial diversity is the Hirschmann Herfindahl Index (HHI) of industry-wise employment. Log(population) is the city's population. Tech Job Share is the ratio of tech job titles as classified by ONET codes over total jobs in a city. Average job tenure is the average number of years workers stay at a job. Job remote suitability is a measure of whether the job-title-by-industry is conducive to remote work (based on job postings of remote work). Standard errors are clustered at the country level.

economic activity with room for productive spillovers, which is associated with larger city effects.

The complexity and diversity of the economic structure allow for linkages between sectors, and network externalities can facilitate higher city income growth. In contrast, a mono-centric industrial structure is associated with lower city effects, suggesting that cities that rely on just a handful of industries are less likely to be productive.

Finally, we test how skill-based agglomeration affects city growth. Previous work argues how knowledge sharing and innovation may result from the higher density of skilled workers (Moretti, 2004). In our data, we measure various aspects of skilled jobs, by looking at their seniority, classifying jobs based on O*NET categories, and determining the share of tech jobs in a location, whether they are suitable for remote work, and whether they are stable. We find that cities with the most skilled, stable, jobs with senior positions also see increased city effect estimates.

Our results highlight that our estimated city effects are strongly associated with multiple measures of city productivity, and are highly robust to the inclusion of more granular country fixed effects. These findings have important implications for policymakers wishing to unpack what makes a city more productive. Cities that attract high-skill workers, a diverse set of industries, and maintain skilled, stable jobs may be likely to see productivity gains.

We also find significant heterogeneity in the strength of effects across a variety of country level indicators. Figure 19 plots similar relationships between city characteristics and city effects, but now split by country-level characteristics. We find evidence that the importance of city size is stronger in developing countries; perhaps, as smaller cities are more unproductive in developing countries compared to smaller cities in developed economies.

Conversely, we find that indicators of skill-based agglomeration are less related to city effects in countries that are poorer, less urbanized, and have lower rates of college attainment. Our findings suggest that density, rather than the industrial composition of a location, is one of the largest drivers of city effects in developing countries around the world.

6.2 Heterogeneity of Variance of City Effects by Country

The global scope of our data provides a unique opportunity to assess whether potential returns to internal migration vary between countries. If the distribution of city effects were uniform across countries, individuals moving from a city in the 25th percentile to one in the 75th percentile within a country would experience comparable wage gains, regardless of the country. However, when the distribution of city effects is narrower in one country (e.g., country X) compared to another (e.g., country Y), the returns to migration would be considerably lower in country X. This highlights the critical role of cross-country heterogeneity in city effect distributions in shaping the potential wage benefits of internal migration and returns to facilitating migration (Bryan and Morten, 2019).

In order to quantify the relative importance of cities in determining wages around the world, we focus on the decomposition results from Section 5.2 that reports the share of the wage difference between above and below median cities that is explained by city effects for each country. We uncover substantial variation in the importance of city effects across countries. Figure 20a plots

Figure 20: Distribution of City Effects

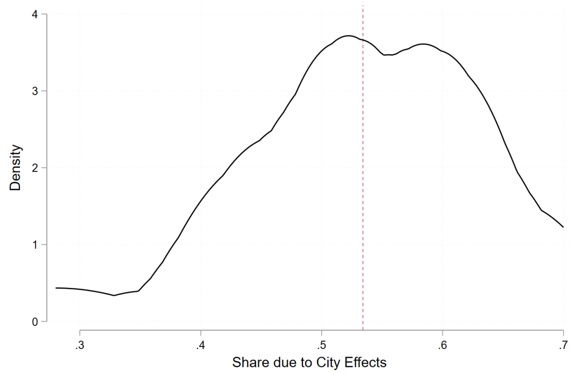

(a) Wage Difference Explained by City Effects

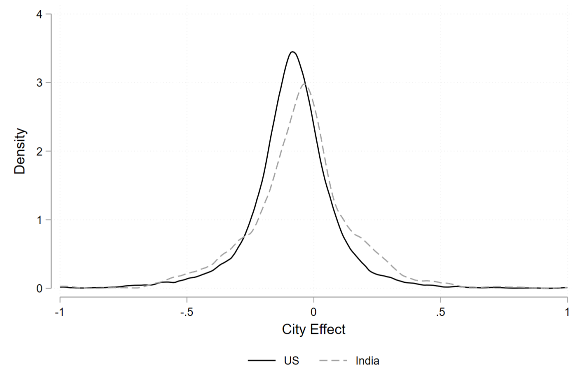

(b) City Effects Distribution, USA and India

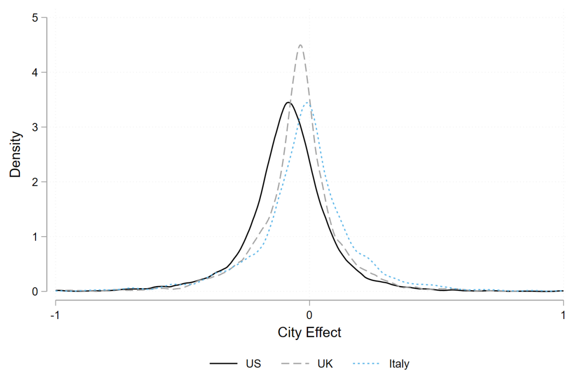

(c) City Effects Distribution, Developed Countries

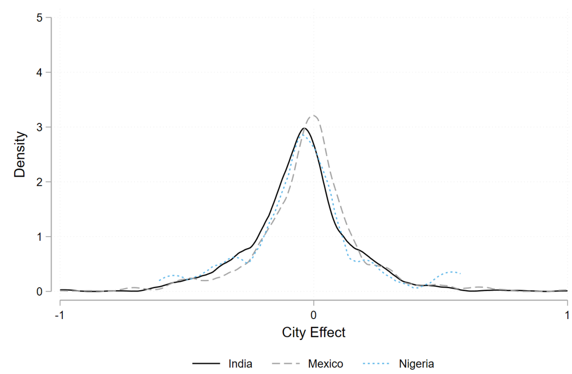

(d) City Effects Distribution, LMICs

Panel (a) plots the density plot of the standard deviation of city effects estimated using Equation 3 across countries. The distribution indicates that the variance of city effects varies meaningfully across countries. Panel (b) shows the distribution of city effects for the USA and India. We recenter the country level distributions, as the city effects for each country are only identified up to a constant term. Panel (a=c) plots the distribution of city effects for a sample of developed countries - USA, UK, and Italy. Panel (d) plots the distribution of city effects for a sample of LMICs - India, Mexico, and Nigeria.

the country-level distribution of the share of wage differences due to city effects. While the average share is large, around 0.53, there is significant global variation in the role of city effects.

In order to highlight the differences in the distribution of city effects across countries, Figures 20b-20d plot the city effect distributions in the US and India, along with a sample of developed and developing countries. Two important patterns immediately emerge; first, the distribution of city effects are different even within developing or developed country categories. Second, generally, the distribution of city effects in developing countries is wider than in developed countries. While this implies that there are more, unproductive cities in developing countries, it also suggests that the potential gains from reallocating individuals across cities are larger.

Indeed, our findings indicate that the potential returns to internal migration differ considerably between countries, depending on the underlying distribution of city effects. For instance, countries with a higher dispersion of city effects–often developing countries–may offer greater economic incentives for individuals to migrate internally, whereas, in countries with lower dispersion, the benefits of such migration might be more constrained.

We evaluate which country-level economic factors are correlated with higher variances of city effects in a country below.

6.3 Explaining the Variance in City Effects

Here, we examine the country-level factors that correlate with the variance of city effects that we estimate. A natural first step is the determine the role that migration plays in shaping the share of wage differences explained by city effects. If migration barriers are high, then individuals are less able to efficiently sort across space. The importance of city effects would thus be higher in such an environment since the role of individual productivity and sorting is limited. Using harmonized, cross-country measures of migration intensity from Bell et al. (2015), we estimate the relationship between migration and the share of wages explained by city effects in a sample of countries in which we have the required data. Figure 21a reveals a strong negative relationship between the internal migration rate of a country and the role of city effects in explaining wage differences. The share of wage differences between above and below-median cities that is explained by city effects is roughly 25% larger in countries with the lowest migration intensities (India and Mexico) compared to those with the highest intensities (Australia and the United States).

Figure 21: Determinants of the Importance of City Effects

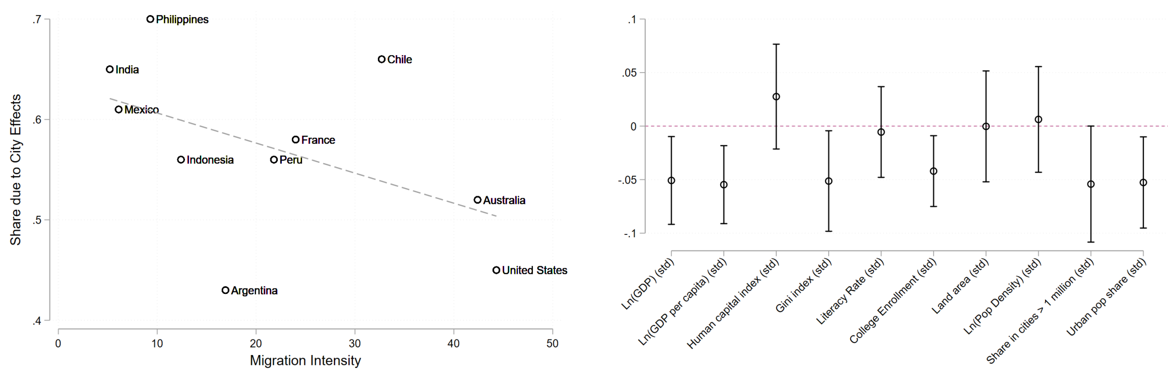

(a) Migration Intensity and Median Share (b) Country Aggregates and Median Share

Panel (a) plots the relationship between the share of wage differences between above and below-median wage cities that is due to differences in city effects within a country along with a line of best fit: $\beta = -.003^{*}(.001)$. Panel (b) plots the coefficients along with 95% CI from regressions on a variety of variables on the same median share as in panel (a). Each variable is standardized, and robust standard errors are implemented.

Additionally, we examine other factors that explain the importance of city effects in a larger sample of countries. All explanatory variables come from the World Bank's World Development Indicators and are calculated as the average value of each variable over the last 10 years.

We look at three important sets of variable categories: income and size, human capital, and urbanization. Figure 21b plots the estimated coefficients for the relationship between the share of wage differences explained by city effects with each of our explanatory variables

Overall, we find that wealthier, more educated, and more urbanized countries have a lower share of wage differences that are explained by city effects. Because a higher share explained due to city effects implies that there would be larger gains from reallocating individuals from low-productivity cities to high-productivity cities within a country, we find evidence that the potential returns to internal migration in LMICs might be higher than in developed countries, consistent with the idea that frictions leading to (spatial) misallocation are higher in developing countries.

7 Conclusion

Our study provides new insights into the economic value of cities by leveraging a unique dataset of detailed job histories from 513 million workers in 220,000 cities across 191 countries. Through an event-study design, we isolate the causal impact of city effects on individual wages, disentangling the role of geographic productivity from ability-based sorting. Our findings underscore the substantial role cities play in shaping earnings, particularly across international borders, where city effects account for up to 93% of observed wage differences. Within-country moves, while also significant, show a more modest contribution of city effects, at around 45-73%.

We find that city effects are strongly associated with population, industrial diversity, city size, and the prevalence of high-skill jobs, reinforcing the importance of agglomeration economies. Notably, the variance in city effects decreases with a country's level of development, suggesting greater potential gains from internal migration in poorer economies. These point to barriers that constrain the realization of these benefits, particularly in developing countries.

Our findings contribute to the broader literature on the economic value of cities, spatial inequality, and migration by illustrating how city-level productivity interacts with national and regional contexts to shape labor market outcomes. Facilitating migration can potentially allow workers to move to more productive cities and raise aggregate incomes. Policymakers aiming to foster economic mobility should consider both reducing barriers to migration and enhancing the productivity of lagging cities. Future research could explore the mechanisms driving these city effects, such as infrastructure investments, governance quality, and social networks, to better inform urban development policies.

References

Abowd, John M., Francis Kramarz, and David N. Margolis (1999) "High wage workers and high wage firms," *Econometrica*, 67 (2), 251–333.

Amanzadeh, Naser, Amir Kermani, and Timothy McQuade (2024) "Return Migration and Human Capital Flows," *Working Paper*.

Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud (2014) "Productive Cities: Sorting, Selection, and Agglomeration," *Journal of Political Economy*, 122 (3), 507–553, http://www.jstor.org/stable/10.1086/675534.

Bell, Martin, Elin Charles-Edwards, Philipp Ueffing, John Stillwell, Marek Kupiszewski, and Dorota Kupiszewska (2015) "Internal Migration and Development: Comparing Migration Intensities Around the World," *Population and Development Review*, 41 (1), 33–58.

Bryan, Gharad and Melanie Morten (2019) "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127 (5), 2229–2268, 10.1086/701807.

Card, David, Jörg Heining, and Patrick Kline (2013) "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128 (3), 967–1015.

Card, David, Jesse Rothstein, and Moises Yi (2023) "Location, location, location," *Working Paper*.

Ciccone, Antonio and Robert E. Hall (1996) "Productivity and the Density of Economic Activity," *The American Economic Review*, 86 (1), 54–70, http://www.jstor.org/stable/2118255.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2008) "Spatial wage disparities: Sorting matters!," *Journal of Urban Economics*, 63 (2), 723–742, https://doi.org/10.1016/j.jue.2007.04.004.

Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Suedekum (2022) "Matching in cities," *Journal of the European Economic Association*, 20 (4), 1478–1521.

Diamond, Rebecca (2016) "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106 (3), 479–524, 10.1257/aer.20131706.

Düben, Christian and Melanie Krause (2021) "Population, light, and the size distribution of cities," *PLOS ONE*, 16 (2), e0245771.

Duranton, Giles and Diego Puga (2004) "Micro-foundations of urban agglomeration economies," in Henderson, J. Vernon and Jacques-François Thisse eds. *Handbook of Urban and Regional Economics*, 4, 2016–2117, Amsterdam: North-Holland.

Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny (2014) "Spatial sorting," *Journal of Political Economy*, 122 (3), 554–620.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams (2016) "Sources of geographic variation in health care: Evidence from patient migration," *The Quarterly Journal of Economics*, 131 (4), 1681–1726.

Glaeser, Edward L. and Joshua D. Gottlieb (2009) "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States," *Journal of Economic Literature*, 47 (4), 983–1028, 10.1257/jel.47.4.983.

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

Gollin, Douglas, David Lagakos, and Michael E. Waugh (2014) "Agricultural Productivity Differences across Countries," *American Economic Review*, 104 (5), 165–170.

Hsieh, Chang-Tai and Peter J. Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124 (4), 1403–1448, 10.1162/qjec.2009.124.4.1403.

de La Roca, Jorge and Diego Puga (2017) "Learning by working in big cities," *The Review of Economic Studies*, 84 (1), 106–142.

Moretti, Enrico (2004) "Human capital externalities in cities," in Henderson, J. Vernon and Jacques-François Thisse eds. *Handbook of Urban and Regional Economics*, 4, 2243–2291, Amsterdam: North-Holland.

Pritchett, Lant (2017) "Lant Pritchett on Poverty, Growth, and Experiments," `https://www.econtalk.org/lant-pritchett-on-poverty-growth-and-experiments/`.

Rosenthal, Stuart S. and William C. Strange (2004) "Evidence on the nature and sources of agglomeration economies," in Henderson, J. Vernon and Jacques-François Thisse eds. *Handbook of Urban and Regional Economics*, 4, 2119–2171, Amsterdam: North-Holland.

Roy, A. D. (1951) "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3 (2), 135–146, `http://www.jstor.org/stable/2662082`.

Young, Alwyn (2013) "Inequality, the Urban-Rural Gap, and Migration," *The Quarterly Journal of Economics*, 128 (4), 1727–1785, 10.1093/qje/qjt025.