The Global Value of Cities

Aakash Bhalothia
UC San Diego

Gavin Engelstad
Northwestern

Gaurav Khanna
UC San Diego*

Harrison Mitchell
UC San Diego

November, 2025
Preliminary

Abstract

We estimate the economic value of cities worldwide using detailed job histories for 513 million workers in 220,000 cities across 191 countries. We use these estimates to document why some cities are more productive, and quantify the earnings gains from migration over the course of development. The data include complete job spells—with start and end dates, establishment names, locations, job titles, and effective salaries—allowing us to implement an event-study movers design with individual and time fixed effects. We find that moving to higher-value cities leads to immediate gains in job seniority, transitions to higher-paying industries and occupations, and large increases in earnings. The global scope of our analysis lets us compare internal and cross-border moves and assess how the productivity advantages of cities vary with development. Across borders, 93% of wage changes are attributable to city effects; within countries, 45–73% are. Richer countries exhibit stronger ability-based sorting, reducing the share explained by location. City effects increase with industrial diversity and population size, consistent with agglomeration economies. More productive cities allocate a greater share of workers to high-productivity firms. The wide dispersion of city effects within countries points to potential gains from internal migration—especially in low-income, less-urbanized economies. Reallocating workers within and across cities to match the US distribution yields substantial wage gains in developing countries.

Keywords: City wage premia, movers design, spatial sorting, gains from migration

*Corresponding author: gakhanna@ucsd.edu
We thank Lyric Greif, Krystal Qiu, and Akshay Yeddanapudi for excellent research assistance. We thank Victor Couture for generously sharing his data. We are grateful to Giovanni Peri, Mushfiq Mobarak, Miguel Zerecero, Vedant Vohra, Nicolas de Roux, various seminar participants, and colleagues for their insightful comments.

1 Introduction

Why are some cities more productive than others, and how much could aggregate incomes rise if more people lived in high-value places? Answering these questions requires measuring the economic value of every city around the world—estimating, for example, the income gains from moving a software engineer from Bangalore to San Francisco. Yet observed wage differences across cities may largely reflect the sorting of more skilled workers into more desirable locations, rather than the causal effect of place itself. Identifying the true impact of location, therefore, requires tracking large numbers of cross-city moves. We leverage a new global database of job histories, covering 513 million high-skilled workers in 220,000 cities worldwide. Using an event-study design that exploits worker movements across cities (Finkelstein et al., 2016; de La Roca and Puga, 2017; Card et al., 2025; Badinski et al., 2023), we first ask: are global disparities in city wages mainly the result of sorting on skill (Young, 2013; Behrens et al., 2014; Combes et al., 2008), or do cities themselves exert a causal influence on worker productivity and pay (Ciccone and Hall, 1996; Diamond, 2016; Glaeser and Gottlieb, 2009; Card et al., 2025)? With our estimated city effects, we then unpack what makes some cities more productive than others, and what the gains are from moving people to high-value cities.

We first distinguish between city effects and individual-level sorting. City effects capture a city's inherent productivity—its firms, industries, technology, infrastructure, and agglomeration forces—while sorting reflects high-ability individuals self-selecting into more desirable locations, generating wage differences driven by worker characteristics rather than place-based productivity. Distinguishing between these mechanisms is essential for understanding how migration affects aggregate income. If city effects dominate, moving workers from low- to high-productivity cities raises aggregate output and yields large migration gains (McKenzie et al., 2010; Clemens, 2011; Clemens et al., 2019).[1] If sorting prevails, cities may instead focus on attracting talent. We estimate city effects for 123,431 cities worldwide using an event-study design that tracks migrants across cities and measures changes in earnings, job seniority, occupation, and industry. This approach quantifies the share of wage differentials between city pairs that translates into actual wage gains for movers. We account for the heterogeneity in premia across firms within a city, by estimating firm effects for all establishments worldwide in our data.

With these estimated city effects, we pursue two main objectives. First, we seek to understand what makes a city more valuable. We estimate establishment effects for all firms worldwide and show that more productive cities allocate a larger share of workers to high-value firms. We relate city effects to observable characteristics—such as economic structure, amenities, industrial composition, and firm diversity—to identify the features that make some cities systematically more productive than others. Our second objective is to quantify how the allocation of workers across firms and locations shapes aggregate productivity. The dispersion of wage premia across cities

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

within a country reflects the potential for higher aggregate income through migration to high-wage locations, and we show that these migration gains evolve over the course of development. In high-income countries, internal migration more frequently directs workers toward productive cities. Our quantification indicates that reallocating workers both *within* and *across* cities could generate substantial aggregate income gains in low-income economies.

We obtain global data on 513 million LinkedIn users, and track their movements across 220,000 cities worldwide. Our data include job location, job titles, firm/establishment names, and the start and end dates of each work spell. We construct individual-level panels using the full job histories of individuals. LinkedIn data captures 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 the availability of precise granular location information, regardless of where a firm is headquartered. This enables us to track the movements of individuals across cities worldwide, along with 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.

With our data, we document 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).

We then highlight descriptive evidence with an event-study analysis of changes in job characteristics with bilateral moves across any two pairs of cities. Conditional on individual, time, and time-since-move fixed effects, we study how moving to a city with high wages, higher fraction of senior positions, high-paying occupations, and high-wage industries, leads to substantial improvements in an individual's earnings, job seniority, occupation, and industry.[2] That is, our event studies show a substantial increase in job quality and earnings when moving to a city with relatively better jobs.

Our event-study design offers several advantages. By conditioning on individual fixed effects, we absorb all time-invariant worker characteristics, and the event-study framework allows us to test for pre-trends in wages prior to a move, verifying that migration is not driven by preceding wage shocks. Identification relies on the assumption that wage shocks at the time of relocation are uncorrelated with individual–city match-specific factors, for which we provide supporting evidence. We then exploit variation across movers to estimate how changes in city location affect wages. Including individual and time-since-move fixed effects allows for systematic mobility correlated with worker and city characteristics—for example, more productive workers may be more mobile, and high-productivity cities may attract such workers. As in the broader literature, our design assumes exogenous mobility and additive separability, implying that relocation decisions are not

[2]The closest in methods is Finkelstein et al. (2016), who study Medicare health outcomes.

driven by unobserved match quality or contemporaneous shocks such as firm closures or rapid human-capital accumulation. Empirically, we find no meaningful pre-trends and show symmetric effects of moving between city pairs, supporting these identifying assumptions.

Yet, simple Abowd et al. (1999) AKM frameworks may yield biased estimates when there is heterogeneity in firm wage premia within cities (Card et al., 2025). For instance, a worker moving from a high-paying firm in a low-wage city to a low-paying firm in a high-wage city would confound the true city effect. To address this, we estimate 'establishment-level' effects for all firms globally and aggregate them to the city level, which serves as our primary measure of city productivity. Additionally, we argue that using imputed wages as a proxy for job quality can bias estimates when worker–city matching is assortative. Under positive sorting, imputed wages may attribute part of the firm effect to individual wage gains. We develop an econometric model that bounds this bias and validate our approach using matched employer–employee data from Italy.[3]

Our estimated city-level productivities show substantial heterogeneity in the share of geographic wage variation explained by these city effects. We show that for international moves, 93% of the wage differentials between any two pairs of cities, translate into gains for the movers. But a lower fraction of other job differences (seniority, occupation, industry) translate into gains. On the other hand, for within-country moves, about 45-73% of the wage differentials translate into gains.[4] 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, with important implications for potential gains from internal migration.

Once we have our estimates of the value of every city around the world, we shift focus to the main objectives of our paper and answer two sets of important questions: First, why are some cities more productive than others? And second, what are the potential aggregate country-level gains from reallocating people *within* cities (across firms), and *across* cities (within a country)?

The path to answering these two questions relies on new important facts that we uncover. First, the allocation of workers across firms, *within* a city, matters. We decompose the city effects into having more high-paying firms, or having a higher share of workers working in relatively high-paying firms. We build on methods developed by Olley and Pakes (1996), and show that the allocation of workers across the firm-productivity distribution is an important component of city effects. That is, there are a higher share of workers concentrated in the most productive firms in high-income countries like the US and UK. In contrast, in lower-income countries (such as India and Mexico), low-productivity cities have too few workers in the most productive firms.

We then relate these estimated city effects to observable city characteristics and uncover several clear patterns. Cities with greater industrial diversity and complexity exhibit larger city effects,

[3] Our method has implications for a wide range of studies relying on imputed data (Abramitzky et al., 2016, 2021; Eli et al., 2016; Amanzadeh et al., 2024). In large-scale datasets such as the U.S. CPS, roughly 30% of wages are imputed, with rates rising over time (Bollinger and Hirsch, 2013).

[4] 50.8% of wage differences in the US reflect city productivity, almost precisely mirroring the 50% finding from Card et al. (2025) on US commuting zones, using the LEHD data.

while more specialized cities with concentrated industry structures have lower city effects. City effects also increase with city size, consistent with agglomeration economies, and with the prevalence of skilled and senior-level jobs (Hornbeck and Keniston, 2017; Greenstone et al., 2010). Together, these patterns both validate our approach to estimating city effects and shed light on a long-standing question in urban economics: what distinguishes more productive cities?

While differences in city effects may reflect potential earnings gains, they may not necessarily reflect welfare improvements. The divergence between measured productivity and welfare may differ across the development spectrum, shaped by policy constraints, congestion, and urban form in developing cities (Akbar et al., 2024; Couture and Handbury, 2020; Harari and Wong, 2025a; Hanna et al., 2017; Harari, 2020). We investigate how amenities are correlated with city effects, and the extent of spatial decay in effects from large cities. We document positive correlations with pollution in developed economies (but not developing), potentially reflecting that productive cities pollute more. Yet, non-temperate climates display compensating differentials. We also document meaningful agglomeration shadows (Hornbeck et al., 2025), whereby proximity to a large city at first helps nearby areas, but at certain distances creates low levels of economic activity.

Our next set of results addresses our second main question: what are the potential aggregate earnings gains from reallocating workers across firms and cities (Albert and Monras, 2022; Desmet and Rossi-Hansberg, 2013; Orefice and Peri, 2025)? We begin by studying gains from moves *within* countries. Potential earnings gains from spatial reallocation rise with migration distance: moves between more distant cities yield larger gains, reflecting greater differences in city effects between origins and destinations. Consistent with this pattern, actual movers tend to relocate over longer distances, potentially capturing higher returns. These patterns, however, vary across income levels. In both rich and poor countries, workers are more likely to move toward higher-wage cities. Yet in developing countries, the highest-wage cities are not always the most productive. As a result, many moves in low-income economies target cities that are not the most productive, implying sizable unrealized gains from more efficient migration patterns.

We next examine the dispersion of city effects within countries, which captures the potential gains from internal migration. This dispersion is substantially larger in low-income, less urbanized countries with fewer major cities. Moreover, the share of wage variation attributable to city effects is higher in countries with lower migration intensity, consistent with migration frictions—particularly in low-income settings—limiting efficient sorting across locations and amplifying the role of place.

Finally, drawing on our estimated city effects and decomposition results, we conduct a few simple counterfactual exercises. We quantify how much average wages could rise if workers were better matched to productive places and firms. First, we reassign each country's population *across* its cities so that the distribution of people over the city-effect spectrum mirrors that of the US. This reveals large, unrealized gains: India's average annual wages rise by 2.3 percent (about 200 USD per person), with even larger improvements in Mexico and Nigeria. Next, using evidence that US cities allocate workers more efficiently across firms within a location, we apply the US pattern

of *within*-city allocative effects to other countries. This boosts average wages by 2.6 percent in India, though the gains are smaller in Mexico, Nigeria, and especially modest in richer countries where such frictions are weaker. Finally, combining both reforms—improving firm-level allocation and then reallocating population across cities—we find that the effects are somewhat additive: India's average wages rise by 4.3 percent, with similar or larger gains elsewhere. Taken together, these results indicate that place-based and firm-level reforms do not crowd each other out and that substantial productivity gains remain locked behind both spatial and within-city frictions.

These findings have important implications for recent debates on the value of locations. A recent starting point is Card et al. (2025), who use AKM methods to isolate gains from locations within the US. The distinctly global nature of our analysis creates certain advantages.[5] 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, we show 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., 2025) or country-level gains (Amanzadeh et al., 2024). Second, our global analysis allows us to understand how the potential gains from reallocating workers across firms and space vary greatly across countries around the world. This reflects recent work on 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) and Badinski et al. (2023)), 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 (Harari, 2020; Asher et al., 2025; Harari and Wong, 2025b; Harari, 2024). For instance, agglomeration forces and skill concentration seem to be more relevant in developing countries. Importantly, in low-income countries, we argue that certain cities have low productivity because few workers reside in the high-productivity firms. Re-allocating workers to higher-productive firms within these cities may improve city productivities.

Third, 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., 2025). 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

[5] See also, other work by Carry et al. (2025); Card et al. (2013); Bonhomme et al. (2019); Borovičková and Shimer (2024); Eeckhout and Kircher (2011).

these are non-negligible city effects, and are correlated meaningful measures of economic structure.

Finally, we speak to the literature on the gains from migration (Clemens et al., 2019; Heise and Porzio, 2021; Clemens, 2011; McKenzie et al., 2010; Khanna et al., 2025). We argue that these potential gains are particularly high in low-income countries, where the dispersion in the importance of city effects is larger. The re-allocation of workers to high-productive urban centers has the potential to raise aggregate incomes in low-income countries.

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, and potential sources of bias. Section 5 discusses the estimated city effects and their implications for heterogeneity in gains from different types of migration. Sections 6-8 documents new facts and implications of these city effects. Section 6 explains what drives these city effects, industrial drivers of productivity, the presence of agglomeration shadows and spatial decay, and importantly, how the allocation of workers across firms contributes to productivity differences. Section 7 documents the distribution of effects and potential gains from migration over the process of economic development. Section 8 quantifies the gains from re-allocation within and across cities, while Section 9 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 focus our discussion on LinkedIn data since that is a large majority of our users. The data contain over 513 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.[6]

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 exhibit varying degrees of LinkedIn use, even among college-educated individuals. As noted in Amanzadeh et al. (2024), the ratio of college-educated individuals to LinkedIn users varies across countries, with those in Central Asia particularly unlikely to use

[6]There is a growing literature using Glassdoor data, for instance, to estimate the wage returns to graduating from certain universities (Martellini et al., 2024).

LinkedIn. Although this could complicate analysis of certain questions, most of our analysis will focus on estimating city wage premia within a single country at a time. We would expect the coverage to vary with income, and in Appendix Figure A1 we plot the relationship between the country coverage rate as measured by Amanzadeh et al. (2024) and GDP per capita and find a moderate positive relationship between the two measures. This is why our empirical results using cross-country comparisons control for the coverage rate (doing so does not affect the estimates).

Seniority, Occupation, and Industry Scores. We first analyze variation in the nature of jobs across space using three outcomes: Job seniority, Occupation scores, and Industry scores.

The Seniority metric is created using an ensemble machine-learning model. First, information about an individual's current job, such as title, company, and industry, is used to generate an initial seniority score. Second, details about an individual's employment history, such as the duration of prior positions and the seniority of those roles, are incorporated to produce a second score. Finally, an individual's age is used to create a third score. These three components are averaged to arrive at a continuous measure of seniority for each individual. To convert this continuous score into a categorical measure, the data provider maps raw scores to the most likely bin using samples of predictions associated with recognizable keywords (e.g., 'junior', 'senior', 'director'). This yields an ordinal classification with seven discrete categories: Entry Level, Junior Level, Associate Level, Manager Level, Director Level, Executive Level, and Senior Executive Level. Each level is associated with illustrative job titles (e.g., Accounting Intern, Attorney, Director of Engineering, CFO).

In addition to seniority, we construct standardized measures of occupation and industry for each employment spell to account for variation in the type of work performed and the economic sector in which it occurs. Occupations are classified using O*NET codes (The Occupational Information Network, 2020), which group job titles into a detailed taxonomy based on required skills, tasks, and work activities. Industries are assigned NAICS codes (U.S. Census Bureau, 2022), which provide a hierarchical classification of firms based on their primary economic activity. For each employment spell, we define an 'Occupation Score' as the average predicted wage among all other observations with the same O*NET code, excluding the focal observation. Similarly, we define an 'Industry Score' as the leave-one-out average wage across all other spells in the same NAICS-defined industry. These scores provide wage-relevant benchmarks that capture the quality of employment.

Wages. In order to measure the wages 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,[7] position-specific information such as tenure and seniority, and company identifiers. Wages are imputed using a regression-based model on a dataset of over 200 million salaries across a wide

[7]Thus, while the wages are expected to capture information relevant to location, there is no mechanical city effect added to the estimated wages.

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).

Locations. 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 can match individuals to specific cities around the world. An advantage of using user-provided location information is that we can specifically identify where individuals work, 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.

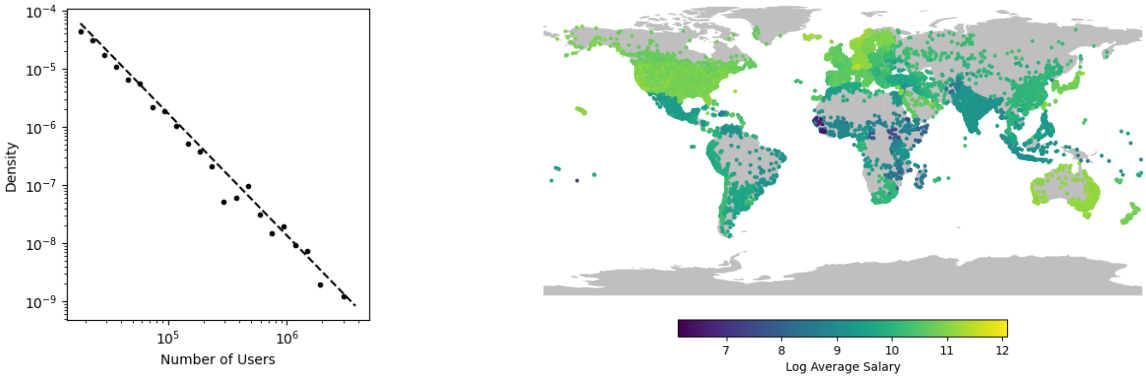
Analysis Sample. We take several steps to construct our primary sample of users. For all work positions with no listed end date–which populate as "present" on LinkedIn–we assign an end date that is equal to the date Revelio Labs collected the data at the end of 2023. We then expand the data from the user-by-position level to user-by-year level by creating yearly entries based on the start and end dates 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. As we show in an extensive set of robustness checks, 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.[8]

3 Descriptive Facts

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).

[8] 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.

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



(a) Cities by Number of Users

(b) Average Log City Wages, Global

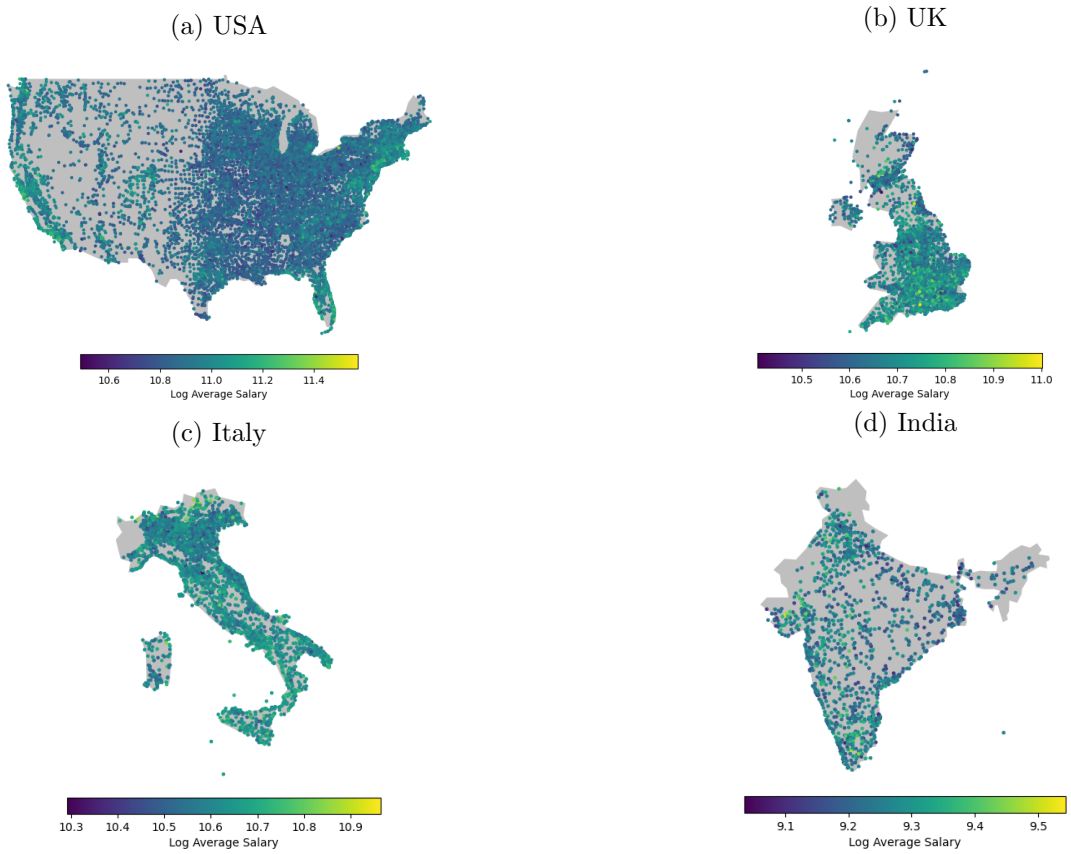
Notes: The left panel shows the density of cities by the number of LinkedIn users. The x-axis is on a log scale and shows the number of LinkedIn users present in our database. The vertical axis is the probability that a city has a certain number of users. The right panel shows the average wages calculated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

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 and the US. 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.

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 US 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.

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.

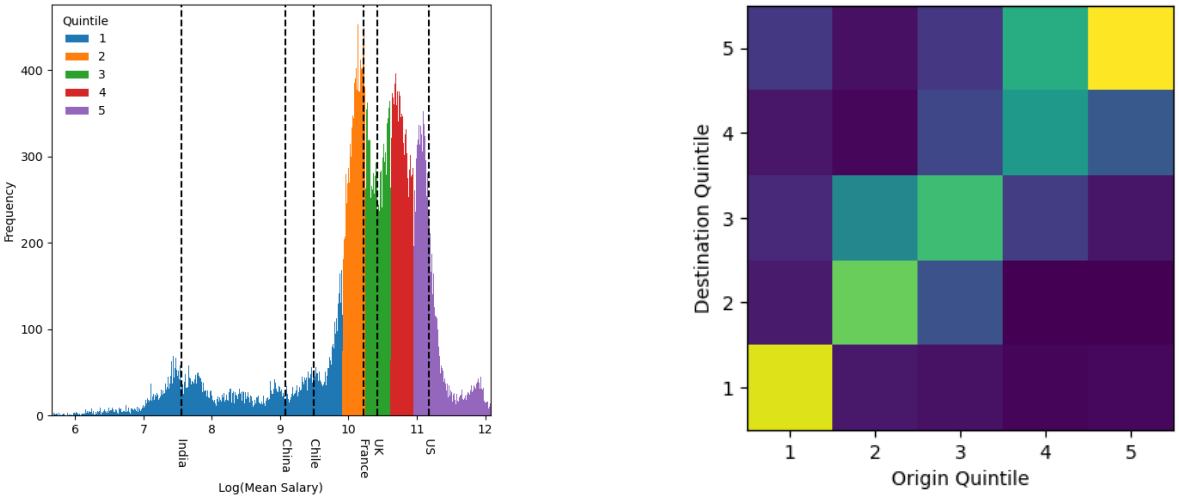
3.2 Transition Matrices and The Movement of Workers Across City Wages

Our estimation will rely on workers who move across cities. We describe this sample of movers, and the moves that they do. Appendix Table A1 compares individuals who move across firms (within cities), across cities (within countries), across countries, and those not in the estimation sample. We might expect the mean salaries to differ across columns for two reasons: First, because different types of people move, and second, because moving itself changes outcomes. While our regressions will control for the former using individual fixed effects, the latter is the primary object of our study. Despite these reasons, the differences across columns are not large, nor statistically different from each other. When looking within countries (like the US or India), even the means across various groups are very similar. This may suggest a similarity between movers and non-movers, but we do not begin our estimation with such an assumption. Our estimation will assume that movers potentially differ from non-movers.

We use our sample of movers to describe not just movers but also moves. Figure 3a shows the distribution of movers by origin income. We divide cities into five quintiles to more easily examine

transitions across location income quintiles. There are more moves seen in the middle region of the city-income distribution. This is consistent with the migration literature, which argues that middle-income regions see the largest out-migration, particularly as low-income individ uals do not have the means for migration, while high-income individuals have less reason to move.

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 five 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.

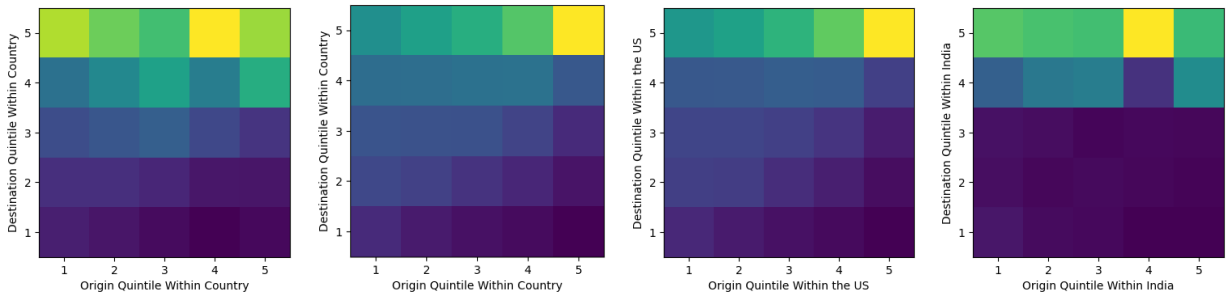
In the next panel, Figure 3b shows the transition matrix of all recorded cross-city moves, with the origin city quintile being on the horizontal axis, and the destination quintile on the vertical axis. Most moves lie along the diagonal. This, perhaps, reflects the fact that most moves are within countries, 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 Figure 4, we examine transitions within and across borders. In Panel A, we study within-country moves. First, higher wage quintiles attract a higher share of internal migrants. This, reassuringly, suggests potential gains to aggregate incomes driven by migration. In low-income countries, there is still some movement to the fourth quintile of cities, but in richer countries, most of the migration occurs to the top quintile, suggesting meaningful allocative gains to income. 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. The top right grid being distinctly higher density suggests, for instance, that those in the top quintile of origin cities move to other cities in the top quintile (and very unlikely to move to cities in the bottom quintile). In India, most individuals move to the top two quintiles, regardless of where they are coming from.

Panel B looks at cross-border moves. The axes are the income quintiles within an origin or destination country.[9] Here, the patterns are even starker. In low-income countries, the top two quintiles are attractive destinations, but in high-income countries, cross-border migrants mostly migrate to the highest-wage cities. This reflects the notion that cross-border migrants may improve allocative efficiency (Orefice and Peri, 2025; Albert and Monras, 2022), despite potential migration frictions (Desmet and Rossi-Hansberg, 2013; Khanna et al., 2025; Hsieh and Moretti, 2019; Bryan and Morten, 2019). Indeed, the US is particularly good at allocating international migrants to high-wage cities, compared to low-income countries.

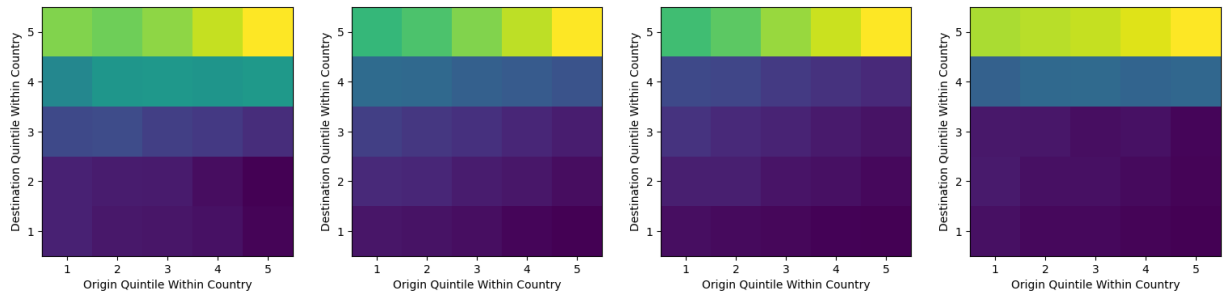
Figure 4: Transition Matrices for Migrants

Panel A: Internal Moves



(a) Low Income Countries (b) High Income Countries (c) In the US (d) In India

Panel B: Cross Border Moves



(e) Into Low Income (f) Into High Income (g) Into the US (h) Into India

Notes: The figures show the transition matrices for within-country moves (top panel) and cross-border moves (bottom 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).

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.

[9] If someone migrates from India to the US, we plot the density for the income quintile of the Indian city, relative to all Indian cities on the origin x-axis, and income quintile of US city, relative to US cities on the destination y-axis.

3.3 Descriptive Event Studies

We begin by 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 wage changes. 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 job characteristics (wages, seniority, occupation score, industry score) 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:

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

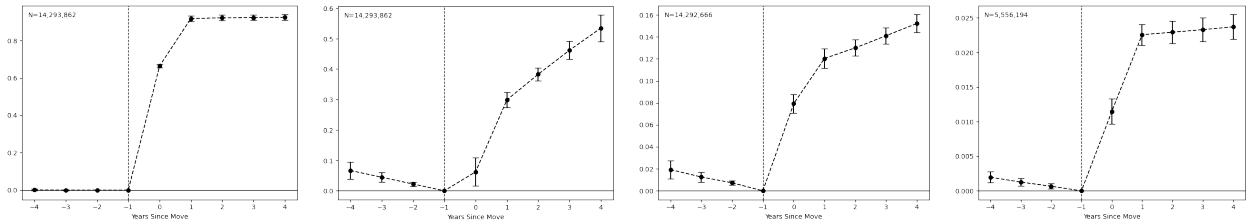
where y_{it} is the outcome (log wage, job seniority, log occupation, industrial score) of individual i in calendar year t. α_i is an individual fixed effect that captures time-invariant skills of worker i, and τ_t controls for calendar year fixed effects. $I_{r(i,t)}$ is a vector for relative-years, 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. For instance, when examining the impacts on earnings, δ_i is the difference in average wages between origin and destination cities, and y_{it} is the wage of the individual i.

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 job outcomes 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 5 Panel A plots estimated coefficients $\theta_{r(i,t)}$ for our entire movers sample. When studying (a) log wages as the outcome, the figure shows a sharp discontinuous jump in the year of the move, from 0 to approximately 0.95, with magnitudes mostly reflecting cross-border moves. This implies a city-effect share of $\approx 95\%$ in the observed variation in wages across cities globally (Finkelstein et al., 2016), and suggests large potential wage gains for international moves, irrespective of individual-level skills. Yet, when examining other outcomes (seniority, occupation score, and industrial score), a lower fraction of the observed variation in wages across cities is because of cities. This suggests the possibility that across international borders, technological differences drive large differences in wages, but not many other job characteristics.

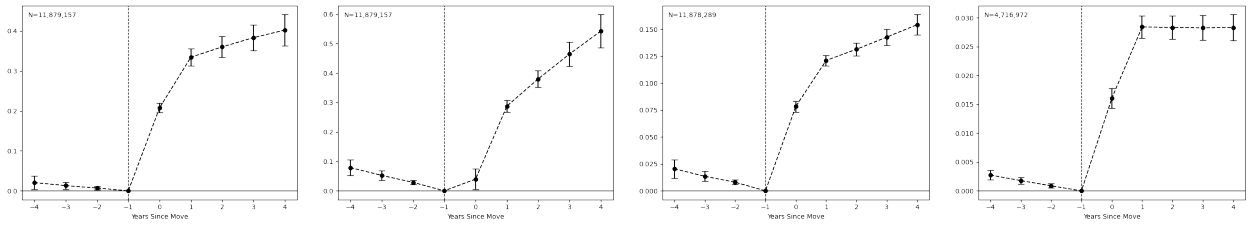
Figure 5: Event Studies by Outcome: Within-Country Movers and All Movers

Panel A: All Movers



(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: Within-Country Movers



(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

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 5 Panel B, shows the results for movers who move to cities within the same country as their origin. For log wages, 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. For non-wage outcomes (seniority, occupation, and industry), the city effects range between 0.08 and 0.5 (encompassing the overall wage effects).

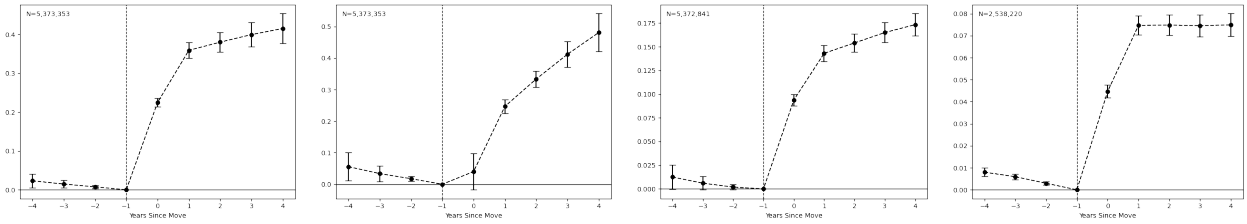
While city effects are meaningful, the much larger wage effects for cross-border moves suggest country effects are extremely important. That is, moving from Bangalore to San Francisco has a much larger impact on one's earnings than moving from Omaha to San Francisco, partly because moving from anywhere in India to anywhere in the US 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. For a few emerging economies, the industry effects are either zero or negative, suggesting that movers to cities with better industries, may actually end up working in worse industries. This may partially reflect frictions in transferring skills across industries, or other job search frictions in low-income settings. Our methods in Section 4.1 allow for this possibility, and our decomposition of city effects (into the allocation of workers across types of firms) confirms this intuition.

We assess the robustness of various decisions made in the data construction in Figures A2 to A4. These figures together show that the patterns we find are mostly unchanged in interpreting

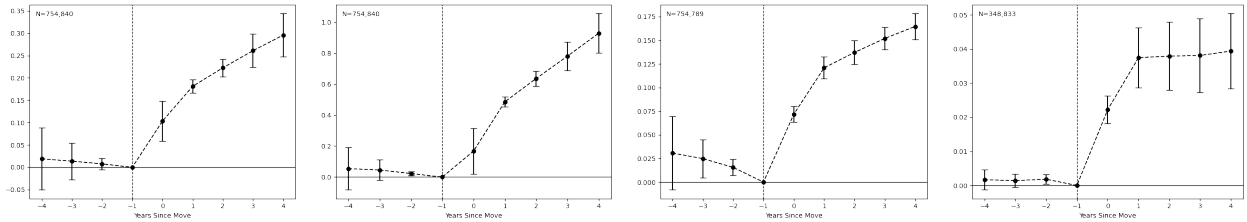
Figure 6: Event Study for Internal Movers, Developed Economies

Panel A: United States



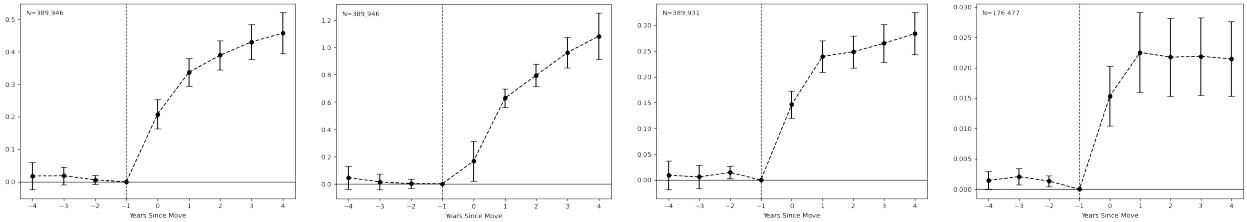
(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: United Kingdom



(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

Panel C: Germany



(i) Log Salary (j) Seniority (k) ONET Wage (l) NAICS Wage

Each panel presents event study coefficients for individuals who moved cities within the respective country. The sample is restricted to individuals employed in the year prior to moving ($t = -1$). Standard errors are clustered at the user level.

the data on salary growth in various ways. For instance, we may think that assuming a constant growth rate for (non-reported) years in between for an individual may affect trends. We show that they do not. These figures cover all the main results, for all countries under focus.

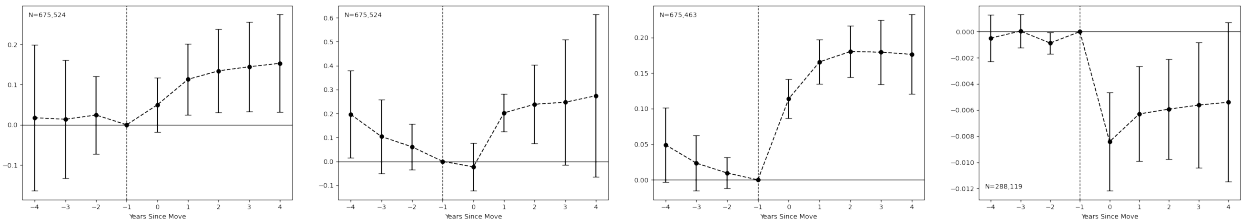
3.4 Average Gains by Origin-Destination Pair

The event study graphs show stark increases in earnings, as a function of the origin-destination wage differential.[10] Keeping with the analysis of bilateral (origin-destination pair) moves, here we describe the relationship between the difference in earnings in the two years prior to and after a move, and the average wage differential at the origin-destination pair level ($\delta_i \equiv \bar{y}_{d(i)} - \bar{y}_{o(i)}$).

[10]While all further exercises can also be done for other job characteristics (seniority, occupation, industry), we succinctly focus on wages as it encompasses all these features.

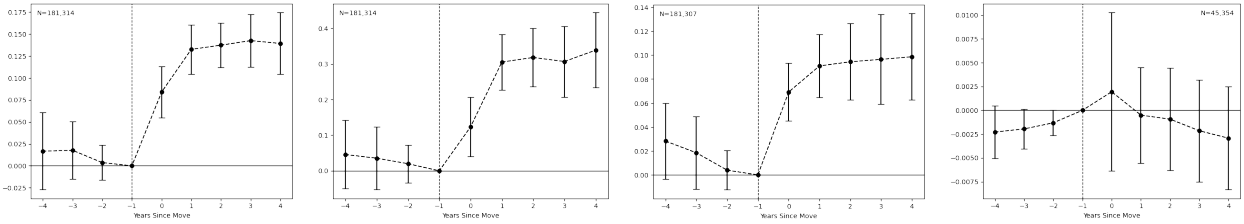
Figure 7: Event Study for Internal Movers, Developing Economies

Panel A: India



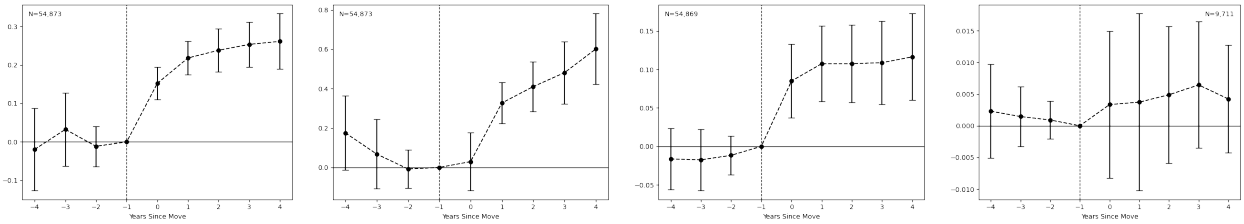
(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: Mexico



(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

Panel C: Chile



(i) Log Salary (j) Seniority (k) ONET Wage (l) NAICS Wage

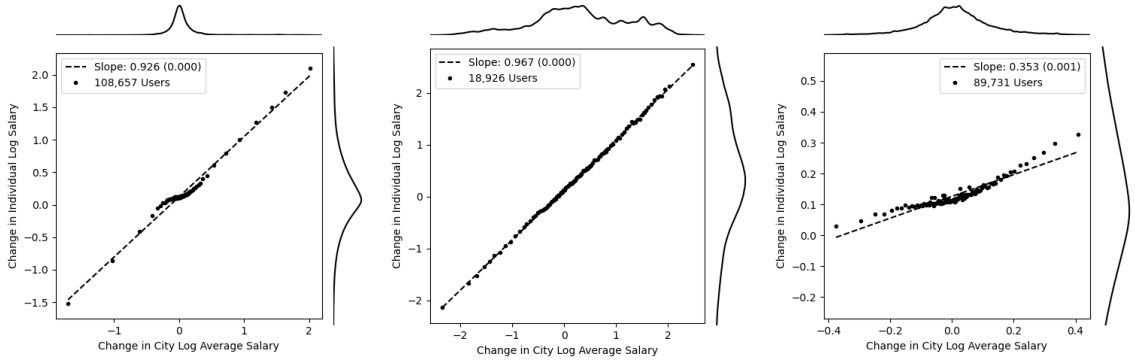
Each panel presents event study coefficients for individuals who moved cities within the respective country. The sample is restricted to individuals employed in the year prior to moving ($t = -1$). Standard errors are clustered at the user level.

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 wage jump 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

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.

reflects that most moves around the center of the graph are for within-country moves with smaller wage differentials, while larger differentials reflect cross-border moves. Since country borders create relatively starker wage differences, moving across countries can lead to potentially larger gains.

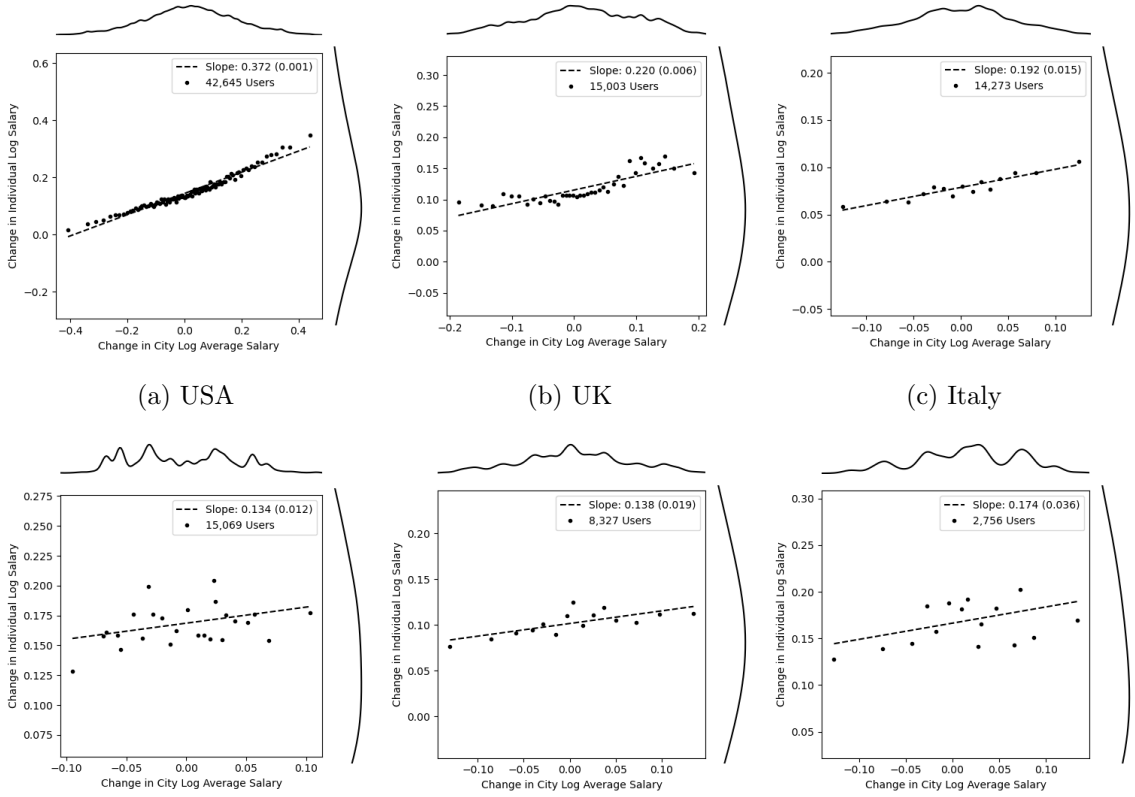
Indeed, the international moves (Figure 8b) show a stark (almost one-to-one) relationship with average city wage differences. The wage differential between Bangalore and San Francisco reflects the fact that, on average, when a worker moves between these cities, they will earn substantially more: almost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

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.

Next, we conduct the same exercise, within specific countries. Figure 9a shows a slope of 0.37 for the US, reflecting a similar size of the event study jump on average. This slope is much flatter for a country like India: Figure 9d shows a slope of 0.13. 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.

We also find that this relationship between wage gains and city wage differentials changes a little over time. Figure A6 plots the same relationship between wage gains over time. While this relationship is constant for international moves, Figure A6b reveals that there has been an increase over time for internal moves; between 2010 and 2020, we find that–for a move between locations

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.

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.

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 the 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. 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 consistent with a simple log-linear production function (so multiplicative in levels) of worker effects and city effects (Eeckhout and Kircher, 2011; Borovičková and Shimer, 2024).

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 temporarily 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. 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 differ from non-movers in levels of worker productivities (with individual fixed effects), and trends in wages around the moves (with time-since-move fixed effects). 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 not mirror the fall for workers who move from B to A.[11]

We test for this symmetry in Figures 8 and 9. We plot 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 high-wage to low-wage cities, 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 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. (2025) refers to as a 'hierarchy effect.'

[11] 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).

Fortunately, Card et al. (2025) provide an intuitive and tractable solution, that we implement. Instead of city-effects, we now estimate establishment effects, and aggregate the establishment 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 were 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 in Section 5.

4.2 Decompositions

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

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

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. Additionally, in the spirit of Card et al. (2025), we can average Equation 4 across workers and years, to decompose differences in mean wages across cities:

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

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

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

We estimate Ω with the help of a simple regression of city effects on average wages ($\Gamma_j = a + \Omega \bar{y}_j + e_j$).

4.3 Expected Bias from Predicted Wages

Our empirical results are based on wage estimates derived from an imputation process conducted by Revelio. 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, position-specific information such as tenure and seniority, and company identifiers. The use of imputed wages to understand the effects of migration is not new.[12] An important potential limitation of this type of data is that the imputation process reduces the variance of wages, in particular because it is unable to fully account for individual-specific determinants of wages.[13] We are unaware of anyone who has formally demonstrated the potential issues of using imputed wages in an AKM framework. We illustrate the potential bias that arises from using imputed wages to estimate city effects, and then propose two solutions to address this issue.

The imputation creates challenges because the prediction model captures firm- and location-level structure more reliably than individual ability. Firm and location are high-signal covariates that recur across many workers, allowing a regression-based imputer to recover those components reasonably well. But individual ability is only proxied by title, occupation, tenure, and seniority, so more of the true α_i ends up in the residual. If higher-ability workers sort into better firms and more productive cities, this residual component of α_i is correlated with city productivity, leading the method to attribute person-driven gains to places instead.

Estimating firm effects first and then aggregating them to the city level somewhat mitigates this problem by netting out the within-city heterogeneity. Movers often come from strong firms in low-wage places and enter weaker firms in high-wage places, so a single city estimate conflates "this city has more high-paying firms" with "migrants do not join the top firms." By assigning wage premia to the exact firms workers enter and only afterward averaging these within cities, the resulting city effect reflects the true firm environment of the city rather than the specific mix of origin and destination firms experienced by movers.

We now consider the implications of using imputed wages that reflect both firm-level characteristics and individual-level observables correlated with worker ability. Suppose the true wage

[12] Examples include Abramitzky et al. (2016), Abramitzky et al. (2021), and Eli et al. (2016) using "wage scores," the average wage in a sector from the 1950 US Census applied backwards in time to full-count census data, or Amanzadeh et al. (2024) using the same LinkedIn data to understand the effects of return migration.

[13] Imputation is common in most large scale surveys and censuses, and rates are high even in gold standard US datasets like the ACS and CPS, which has 30% earnings imputation rates (Bollinger and Hirsch, 2006, 2013).

structure follows the standard two-way fixed effects model:[14]

$$y_{it} = \alpha_i + \gamma_{\boldsymbol{f}(i,t)} + \varepsilon_{it}, \tag{9}$$

where α_i is the worker fixed effect, $\gamma_{\boldsymbol{f}(i,t)}$ is the firm effect, and ε_{it} is an idiosyncratic shock. In our setting, we do not observe y_{it} directly, but instead observe an *imputed* wage $\hat{y}_{it}$ that depends in some way on the average wage at the firm and individual-level covariates such as tenure, education, or job title:

$$\hat{y}_{it} = f(\alpha_i, \gamma_{\boldsymbol{f}(i,t)}) + \xi_{it} = \rho_{\alpha i}\alpha_i + \rho_{\gamma i}\gamma_{\boldsymbol{f}(i,t)} + \xi_{it}, \tag{10}$$

where f is some imputation function that predicts wages based on individual and firm-specific information. In order to more clearly see the potential issues with this approach, for the sake of expositional clarity, we assume the imputation process takes a linear form. Here $\rho_{\alpha i}$ and $\rho_{\gamma i}$ reflect the extent to which the imputation process captures the true worker and firm effects, respectively, and ξ_{it} is a mean-zero residual orthogonal to both α_i and $\gamma_{\boldsymbol{f}(i,t)}$. Substituting the true wage into the AKM regression, we have:

$$\hat{y}_{it} = \alpha_i + \gamma_{\boldsymbol{f}(i,t)} + \underbrace{(\rho_{\alpha i} - 1)\alpha_i + (\rho_{\gamma i} - 1)\gamma_{\boldsymbol{f}(i,t)} + \xi_{it}}_{\equiv \eta_{it}}. \tag{11}$$

As η_{it} contains components of both the worker and firm effects, it is mechanically correlated with the regressors in the AKM specification. This correlation arises because the imputed wage $\hat{y}_{it}$ reflects only partial information about the true fixed effects, capturing some share of the variation in α_i and γ_f, but not all of it. So, when we estimate the AKM model using $\hat{y}_{it}$ in place of the true wage, the residual includes structured error terms that are correlated with both α_i and γ_f.

If there is positive assortative matching between workers and firms, then the regression will misattribute part of the variation in worker ability to the firm fixed effect.[15] The direction and severity of the resulting bias depend on the relative accuracy of the imputation. If firm effects are well approximated ($\rho_{\gamma i} \approx 1$) while worker effects are only partially captured ($\rho_{\alpha i} < 1$), then estimated firm effects will be systematically upward-biased. We test this pattern in our empirical analysis using observed wage data.

A direct consequence of this bias is that the variance of the estimated firm effects will also be overstated. Since the regression mistakenly attributes some of the missing variation in worker ability to firms, the dispersion of firm effects estimated using imputed wages will be inflated relative to what would be found using true wages. In Appendix B we construct a conservative lower bound on the share of city-level wage variation explained by true city effects based on the wage model proposed in Equation 9 which we then apply to our results in the subsequent section.

[14] We suppress time fixed effects and additional controls for clarity.

[15] In the absence of assortative matching, we may still be concerned about imputation 'measurement error.'

4.3.1 Empirical Estimates of Imputation Bounds

In addition to our conservative bounding exercise, we next turn to quantifying the potential excess variance that we are attributing to city effects using matched employer-employee data from Italy.[16] Using data covering the work histories of Italians between 1975-2001 and including firm identifiers, location information, tenure at the firm, and a course measure of salary type (to proxy for an occupation code), we are then able to estimate the true city effects alongside our own "imputed city effects." We construct imputed wages for full time workers by taking the average wage for individuals with the same firm, tenure class, salary type, and year; based on documentation of the imputation process provided by Revelio, we are leveraging similar information through this process. We obtain realistic estimates of the bias from using imputed wages by estimating city effects using Italian data, where we observe the true wages of individuals. This allows us to estimate true city effects as well as conduct our own wage imputation and separately estimate an "imputed city effect."

We find that bias from using imputed wages is small and uncorrelated with observable city characteristics. Using these bias estimates, we adjust downward the share of variance in average city wages attributed to city effects. We also verify that our lower-bound approach indeed understates the importance of city effects. Importantly, we find that the resulting lower bound is close to the true share of variance in actual wages explained by city effects.

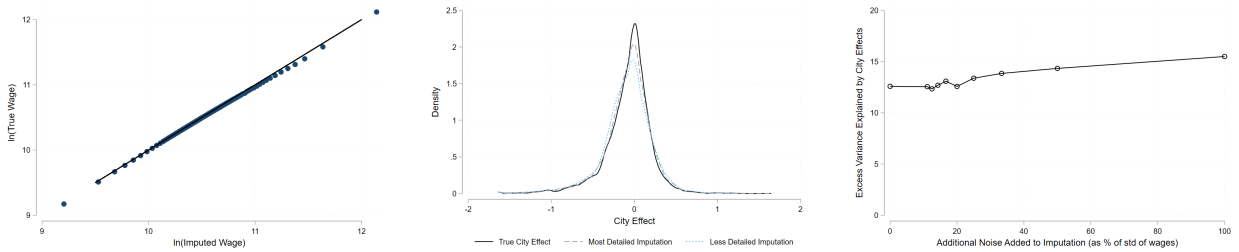
Our first finding is that an imputation process that included all the information above performs very well. A regression of imputed wages on real wages yields an R^2 of 0.875, meaning there is relatively little variation in individual wages once accounting for firm, time, tenure, and salary type. We also analyze imputed wages only using information on firm and salary type, where the corresponding R^2 is smaller at 0.766. Figure 10a plots the relationship between our most detailed imputed wages and real wages; we find a tight relationship between the two measures. Using both real and imputed wages, we estimate the city effects according to the process described above.

Figure 10b plots the distribution of estimated city effects using true wages along with both of our imputed wages. Confirming our prior, we find that the variance in estimated city effects themselves is more dispersed for imputed wages than the true values. Finally, we compute the share of variance in (imputed) wages explained by (imputed) city effects. We find that we overestimate the role of city effects in explaining wage variation by 12.5% (7.3 pp). Importantly, we find no evidence that the excess variance explained by city effects is correlated with city-level observables such as the number of firms or workers in a location or the average true wage of a city. Therefore, we do not find any evidence to suggest that our estimated city effects using our LinkedIn sample are *systematically* biased in any way. However, the excess variance explained by city effects that we estimate diverges from the true value when the imputation itself is inaccurate, especially when individual-specific information, such as tenure, is omitted.

An additional potential concern with using imputed wages is that the imputation accuracy

[16] Dataset developed by Giuseppe Tattara and the Economics Department in Universita Ca' Foscari Venezia.

Figure 10: Bias in city effect estimation arising from wage imputation



(a) Binned scatter of true vs. imputed wages

(b) Distribution of city effects under different imputations

(c) Excess variance explained by city effects under added noise

Panel (a) plots the relationship between imputed wages and true wages using a binned scatter plot, along with a 45-degree line. Panel (b) plots the distribution of estimated city effects when using real and imputed wages, where the two distributions represent different imputation processes. Panel (c) displays the estimated bias in the share of the variance in wages explained by city effects as we add additional noise to the imputation process for wages. Mean 0 noise with standard deviation equal to a percentage of the standard deviation of wages is displayed on the x-axis.

could vary across countries. Perhaps in developing countries, the imputation process is noisier due to sparser information on firms or the effects of tenure. In theory, this could lead us to overstate the importance of city effects in countries with less precise imputation performance. To assess if this poses practical problems, we progressively add random noise to our imputed wages and recompute the share of variance in our imputed wages that is explained by city effects. We define our additional noise variance as a percentage of the standard deviation of the imputed wage. We find little evidence that this type of noise significantly affects our results. Figure 10c plots the relationship between the excess variance due to city effects and the noise added. For values of noise less than 1 standard deviation of the imputed wage, we find an extremely weak relationship between noise and excess variance; adding noise with variance equal to 1 standard deviation of the imputed wage only increases the estimated importance of city effects by roughly 3pp.

Finally, we perform our bounding exercise as described and defined in Appendix B. We find that the lower bound for the share of variance in wages explained by city effects underestimates the effects of cities by 5.7% and 9.0% for the most detailed and less detailed imputation, respectively. Despite the fact that the imputation process conducted by Revelio uses an even richer set of information than we have in our application with Italian data, we will perform our bounding exercise with a conservative estimate of the explanatory power of imputed wages on real wages of 0.875.

4.4 Does our Estimation Truly Capture City Effects?

Our focus is not on methods, but rather our contribution is on the global nature of our analysis, allowing us to uncover meaningful facts about how cities vary across the process of economic development. While we build on established methods (Card et al., 2025; Finkelstein et al., 2016), we explore the intricacies of our estimation strategy here. While we have already extensively discussed the exogenous mobility assumption, there may be other concerns. We do not have definitive answers

to all concerns, but we try to be transparent on possible violations of these assumptions.

First, we are identifying effects off of movers, and they may differ in important, unobserved ways, from those who stay. These differences may be of various kinds. For instance, those with the lowest cost of moving cities, or unobservably talented, may be more likely to move. Table A1 suggests that movers and non-movers are similar in many respects. And differences in average ability are controlled for with the help of individual fixed effects.

However, even after accounting for average differences, there may be systematic differences in 'returns' to moving cities. Movers are likely to benefit more from moves, especially as they are likely to move to more productive cities. That is, people with higher gains are more likely to move cities. These may not necessarily be sources of bias (as the AKM literature highlights), but rather an indication that estimated city effects are 'local' average effects for movers who choose these cities. Yet, a few aspects of our estimation help. The symmetry in effects seen in Figures 8 and 9 suggest that the effects are similar when moving from Omaha to New York, as is vice versa. Furthermore, we highlight that much of our estimates rest on people moving across firms *within* Omaha, and within New York, rather than across cities. Finally, as we later show in Figures 14 and A8, our residuals are flat with respect to city average wages, suggesting there are no other unobserved factors that violate the exogenous mobility assumption.

Another set of concerns could be that even after accounting for permanent differences across workers (with say, individual fixed effects), there could still be unobserved, time-varying factors that influence both the decision to move, and the resulting post-move wage trajectory. Other sets of 'dynamic' possibilities are explored in Bonhomme et al. (2019); de La Roca and Puga (2017). While these are critiques of the variance decomposition exercise, rather than of estimating mean city effects, all our estimates include time-since-moved fixed effects, and our event studies extensively investigate pre-trends, or post-move spikes. We document the entire trajectory of pre-move trends, and post-move dynamics. For instance, if workers experience positive productivity shocks, acquire new skills, become more motivated, and use those opportunities to move to a higher-paying firm, then we would expect to see pre-trends in the event studies.

A third set of issues explores selection into the sample. The coverage rate of LinkedIn data varies by country, in likely systematic ways. People in certain locations may be more likely to report their 'successes' on LinkedIn, and less likely to report 'failures.' Different types of people select into LinkedIn in San Francisco and Lagos. While these are not concerns for estimates of city effects using within-city moves, they are likely to matter when we later investigate how city effects vary across countries and economic development (if the selection into the sample is systematic). This is why, in later analyses, we control for coverage rates extensively, and find that they do not affect our estimates. Indeed, although our coverage rate increases substantially over time (especially post-2006), our event-study coefficients remain relatively stable. A related issue is that developing countries have worse data quality. Our discussion of Figure 10c allows us to bound effects by data quality, and we estimate fairly tight bounds.

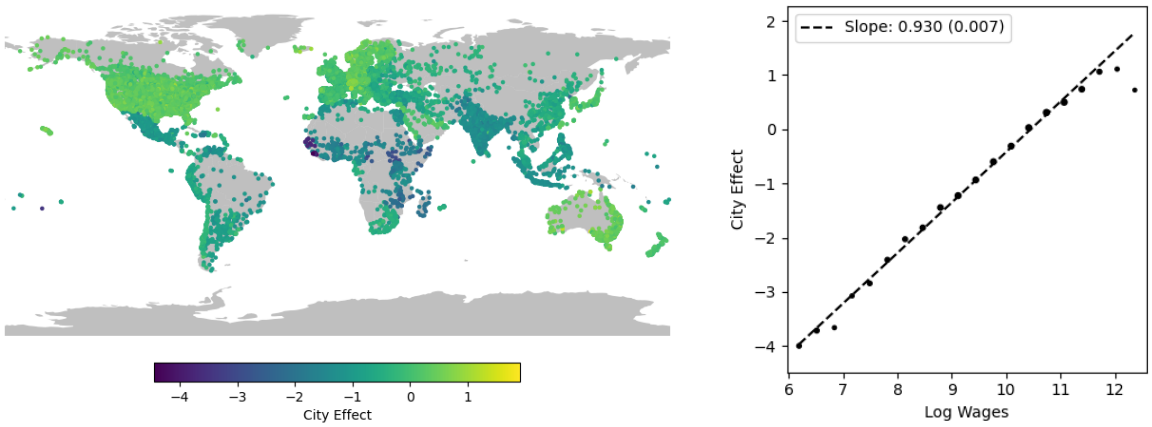
A fourth set surrounds what wages capture (employer market power, prices, etc.). For instance, movers to a city may have less 'bargaining' power (if firms have market power) and thus earn less than their marginal product. In some ways, using wages 'predicted' by seniority/job title/firm are somewhat advantageous. Furthermore, the fact that much of our estimates rest on within-city moves reduces the likelihood that we are confounded by this. Similarly, if wages reflect price differences, relying on within-city moves again helps.

Furthermore, it is important to consider selection into certain pairs of moves. Those who move from Houston to Nairobi may be more likely to be say, oil executives, or government officials, and may earn higher salaries. This is precisely the set of issues the establishment-effect analysis aims to solve. Intuitively, the identification comes from those moving within these cities, across establishments. Yet, even in the city-level event studies, we do not see any such phenomena: that is, we see positive correlations between average wages, and city effects, and that moves from Houston to Nairobi have symmetric, similar (opposite) effects as those from Nairobi to Houston.

A final set of issues concerns limited mobility bias (Kline, 2024), relying on cities with very few moves (Bonhomme et al., 2019; Kline et al., 2020). This is exacerbated when identifying second moments (as in traditional AKM decompositions). We follow the recommendations in Bonhomme et al. (2019) by relying on cities and firms with a substantial number of moves, and maintain our focus on first moments (the mean city effects). Indeed, Card et al. (2025) argues that even the variance decompositions are unbiased when aggregating establishment effects to the city level.

5 City Effects

Figure 11: City Effects, Global

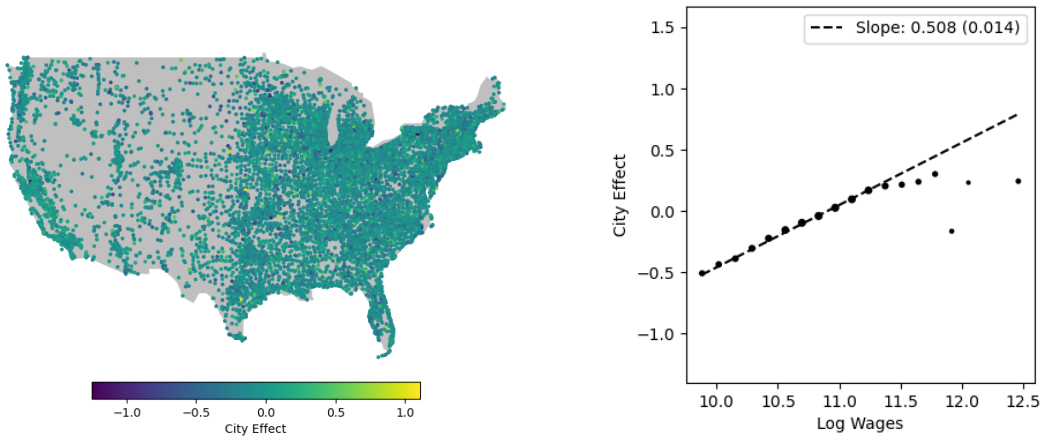


(a) Estimated City Effect

(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 5. 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. $N = 8,025,420$

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. 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.

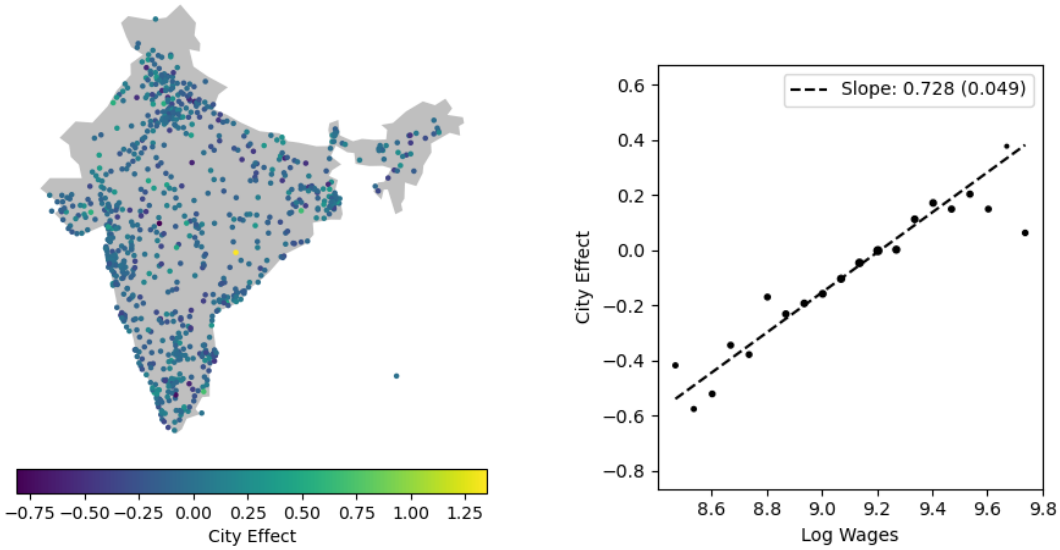
Using our movers' sample and Equations 4 and 5, we estimate city effects for cities worldwide. To assess the economic implications of these estimates, we conduct several complementary analyses. First, we plot city effects for all cities using international movers (Figure 11) and compare them with estimates based on within-country movers (Figures 12, 13, and A7). This comparison illustrates how national borders shape measured city effects. Taken together, these analyses allow us to assess the role of city effects under different migration costs: international moves across countries, moves across cities within a country, and within-country moves in different national contexts. In subsequent sections, we examine the underlying economic primitives captured by these city effects and quantify the potential gains from migration within and across countries.

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 worldwide, highlighting the importance of city-specific factors in shaping economic outcomes.

Figures 12 and 13 focus on within-country moves in the United States and India. The resulting maps differ notably from the raw salary patterns shown in Figure 2. In the US, raw wages are substantially higher on the coasts than in the Midwest, whereas the estimated city effects in Figure 12 are more evenly distributed across space. The flatter relationship between city effects (y-axis) and

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. 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.

average city wages (x-axis) for the US ($\beta = 0.51$) and India ($\beta = 0.73$), relative to the international comparison ($\beta = 0.93$), indicates that city effects account for a smaller share of within-country wage dispersion than of cross-country wage differences.

These findings align with the event-study evidence in Figure 5, which shows that potential wage gains from international migration substantially exceed those from internal migration. Moreover, the heterogeneity in slopes across countries suggests that the role of city effects in explaining wage differentials varies meaningfully by national context. As a result, the potential wage gains from internal migration differ across countries, with larger gains in settings where city effects contribute more strongly to wage determination. We examine the cross-country distribution of city effects systematically in subsequent sections.

5.2 The Contribution of City Effects to Wages

Table 1 reports a decomposition of cross-city wage variation worldwide, estimated using Equation 6. Tables 2 and 3 present analogous decompositions for developed countries (the US, UK, and Italy) and developing countries (India, Mexico, and Nigeria), respectively. Column (1) compares cities with above- and below-median wages, Column (2) contrasts cities in the top and bottom quartiles, and Column (3) compares the top and bottom 5 percent of the wage distribution. Column (4) reports differences between selected city pairs. The final row of Column (5) reports the share of

variance in mean city wages attributable to city effects, Ω, as defined in Equation 8. This statistic is equivalent to the slopes shown in Figures 11, 12, and 13.

Table 1: Decomposition of Wage Differences: Globally

	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
Bounded Share					(0.75)

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). Column (5) reports the overall variance in imputed wages explained by city effects, while the number in the parenthesis is the bound for the explained share based on Equation 26.

Table 1 shows that 93% of the global variation in wages across cities is due to city effects. We compute the bound from Equation 26 and report the value in the parentheses in Column 5, estimated to be 75%. 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, 45–51% of the variation in wages across cities is attributable to city effects (Table 2). Allowing for potential bias from wage imputation, Equation 26 implies a more conservative range of 38–43%. Differences in city effects account for 45–57% of the wage gap between below- and above-median cities in developed countries. Focusing on the US, city effects explain approximately 51% of cross-city wage variation, closely in line with the 50% reported by Card et al. (2025) for commuting zones. Similarly, for the gap between below- and above-median-wage cities, differences in city effects explain about 45% of the variation.

For our sample of developing countries, city effects account for 48–73% of wage differences, with a conservative bound of 40–61% under potential imputation bias (Table 3). Overall, city effects explain a substantial share of wage variation across cities and across different partitions of the wage distribution. Moreover, the share of wage variation explained by city effects is systematically higher in developing countries—a pattern also seen in Figure A7, and which we test formally in Section 6.

In Appendix Table A2 we show the entire variance decomposition. The signs and magnitudes are similar to what Card et al. (2025) find, with some small nuances. While the covariance between user and firm effects is positive in the US, they are negative in certain countries like India. These reflect possible differences in assortative matching and efficiency, we discuss further in Section 6.1.

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
Bounded Share					(0.41)

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
Bounded Share					(0.4)

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
Bounded Share					(0.37)

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). Column (5) reports the overall variance in imputed wages explained by city effects, while the number in the parenthesis is the bound for the explained share based on Equation 26.

5.3 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 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. (2025).

Panel A of Figure 14 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 calcu-

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
Bounded Share					(0.59)

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
Bounded Share					(0.43)

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
Bounded Share					(0.39)

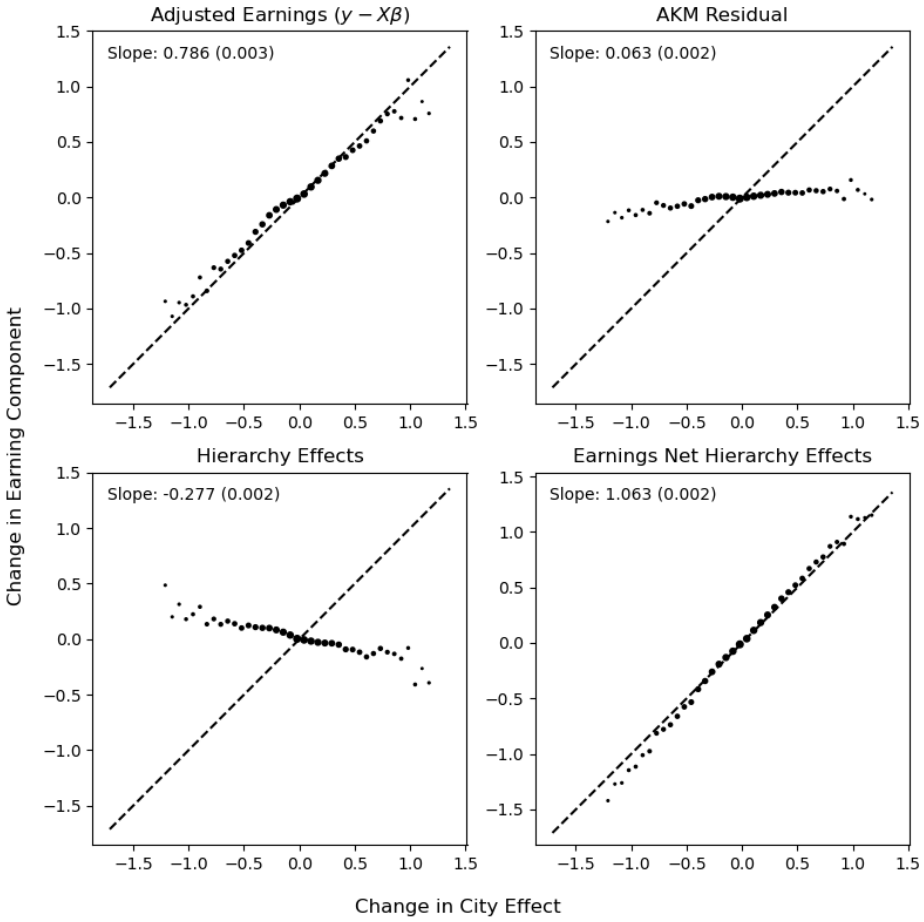
Notes: This table presents the decomposition of wage differences across cities in different developing countries. In each panel, the first row computes the difference in average wages between the destination and origin region. The second row reports the difference in average city effects between the 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). Column (5) reports the overall variance in imputed wages explained by city effects, while the number in the parenthesis is the upper bound for the explained share based on Equation 26.

lated 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 experience 78.6% of the change in earnings we would predict based on changes in city effects when they move.

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

$$y_{it} - y_{i,t-1} = \left(\Gamma_{c(i,t))} - \Gamma_{c(i,t-1)}\right) + \left(h_{f(i,t)} - h_{f(i,t-1)}\right) + \left(\epsilon_{it} - \epsilon_{it-1}\right) \ , \tag{12}$$

Figure 14: 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. We plot a 45' line as reference. The sample is restricted to movers who move once and within their countries.

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 changes in city effects, it would suggest a violation of the exogenous mobility assumption. We test this in Panel B of Figure 14. 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 14. 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. (2025), 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 A8) 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.

6 What Explains the City Effects and Their Distribution

6.1 The Allocation of Workers Across Firms

City effects may partially depend on the fact that productive firms locate in certain locations Carry et al. (2025), and/or on workers being more or less likely to work in highly productive firms. We first decompose our estimated city effects to determine how much of the variation in city effects is driven by employment allocation across high- and low-productive firms. Based on our definition of a city effect, we can build on methods developed by Olley and Pakes (1996) to rewrite our estimate in the following way:

$$\Gamma_j = \sum_{j(f)=j} s_f \gamma_f = \bar{\gamma}_j + \sum_{j(f)=j} (s_f - \bar{s}_j)(\gamma_f - \bar{\gamma}_j) \,, \tag{13}$$

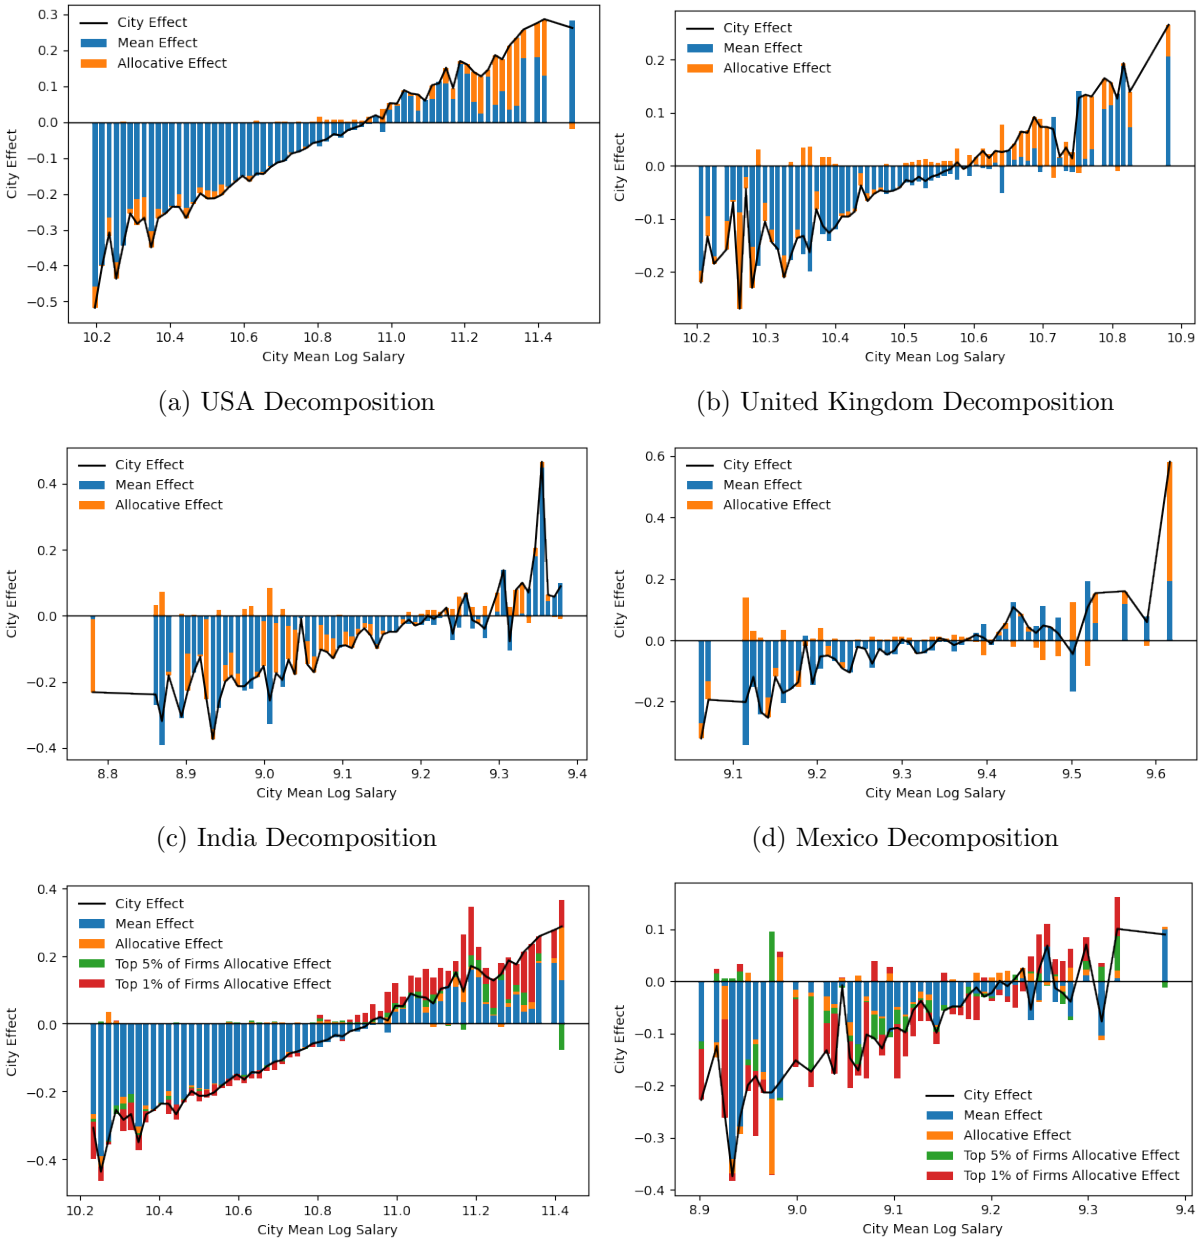
where $\bar{\gamma}_j$ is the average firm effect within a city and $\bar{s}_j$ is the average employment share within the city. We can then estimate how much of the estimated city effect comes from the average firm effect within the city and how much comes from the allocation of workers across firms.

If the above-average firms are also paired with more workers within a city, we would find larger positive city effects. Figure 15a- 15d plots the results from this decomposition across the distribution of average city wages within country, while Figure 15e and 15f further splits the decomposition to separately account for the role the largest firms within a country play.

In the US and UK, we find strong evidence that the allocation of workers across the firm-productivity distribution is a crucial component of the overall city effect; the most productive cities also have more workers employed in the most productive firms, whereas the least productive cities have too few workers in the most productive firms. However, in India and Mexico, fewer workers are allocated to the most productive firms. In the least productive cities, the most productive firms do not employ enough workers, while in the most productive cities, we see weak or no evidence that the most productive firms also hire more labor.

This evidence is consistent with other studies that argue misallocation to be a more significant impediment to overall growth in developing countries. Here, we suggest that the allocation of workers *within* a city lowers overall city productivities in developing countries. When we split the share effect into the effect coming from the largest firms, we find strong evidence that the largest, "superstar" firms within the US are driving a large fraction of the share effect, meaning that in the US, the largest firms are also among the most productive. On the other hand, in India, we find

Figure 15: Decomposition of City Effects According to Employment Allocation



(a) USA Decomposition

(b) United Kingdom Decomposition

(c) India Decomposition

(d) Mexico Decomposition

(e) USA Decomposition, with Superstar Firms

(f) India Decomposition, with Superstar Firms

Notes: We first create 75 bins of cities for their mean log salary. Then, we plot the average Olley and Pakes (1996) components within each bin using Equation (13). The components split the city effect into the mean part (what the city effect would be if all firms had equal shares and effects) and the allocation component (the part of the city effect coming from firm heterogeneity). The sum of the two components is the city effect. Panels (e) and (d) separate the allocation component into the largest 1% of firms, the largest 5% of firms, and the remaining firms.

that the largest firms are often less productive than the average firm within a city.

6.2 Explaining the City Effects

We next examine whether our estimated city effects are associated with observable factors that we believe would affect a city's impact on earnings. Using the full sample of LinkedIn users, we construct a measure of a city's industrial diversity based on employee job classifications, the ratio of job titles identified as senior positions, the average job tenure, and the average remote suitability of positions. Additionally, we obtain a measure of each city's current population by merging our cities with publicly available GeoNames data. We estimate the relationship between city effects and each of these variables, and include region fixed effects to control for heterogeneity across regions.[17]

Figure 16a reports standardized coefficients from our estimates, while Figures 16b–16f present binned scatter plots for each relationship. Consistent with agglomeration theories, we find that larger cities and those with modern firms—proxied by remote-work suitability—exhibit higher city effects. These patterns align with work emphasizing how density and scale facilitate knowledge spillovers that raise urban incomes (Duranton and Puga, 2004; Rosenthal and Strange, 2004).

Next, Figure 16b shows that cities with lower industrial HHI exhibit larger city effects. More complex and diversified economic structures support sectoral linkages and network externalities that raise productivity and incomes. In contrast, a mono-centric industrial structure is associated with lower city effects: cities relying on just a handful of industries are less productive.

Finally, we examine the role of skill-based agglomeration in shaping city effects. Prior work argues that higher densities of skilled workers foster innovation (Moretti, 2004). We capture skill intensity using multiple measures, including job seniority, O*NET-based occupational classifications, the share of tech jobs, remote-work suitability, and job stability. Across these dimensions, cities with more skilled, stable, and senior job structures exhibit higher estimated city effects.

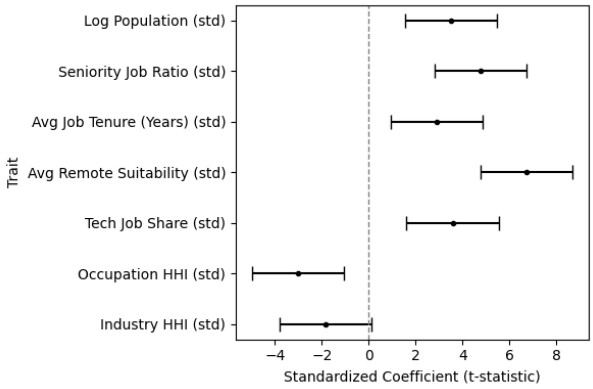
We also find significant heterogeneity in the strength of effects across a variety of country-level indicators. Figure 17 plots similar relationships between city characteristics and city effects, but now split by country-level characteristics. We find evidence that the importance of city size is stronger in developing countries; perhaps, as smaller cities are more unproductive in developing countries compared to smaller cities in developed economies.

We further explore the correlation with each industry in Appendix Figure A9. Along with computer science occupations, other high-skill professional occupations, such as management and legal professions, are positively correlated with higher city effects. Occupations in construction and extractive industries also reflect higher city effects. On the other hand, occupations in retail, personal care, and food services are negatively correlated with city effects.

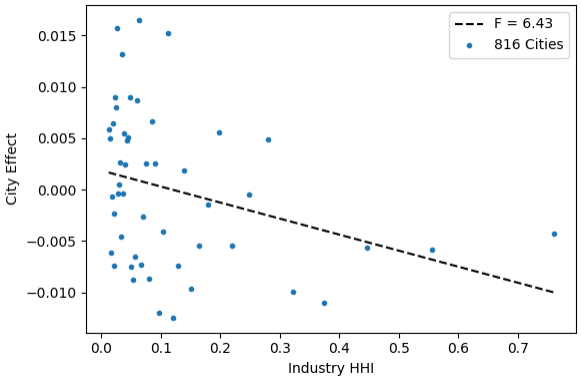
Our city effects are strongly associated with measures of city productivity, and are robust to the inclusion of more granular fixed effects. These findings have important implications for policymakers wishing to unpack what makes cities productive. Cities that attract high-skill workers, a diverse

[17] We identify 15 regions in our data: Arab States, Central America and the Caribbean, Central and Western Asia, Eastern Asia, Eastern Europe, Northern Africa, Northern America, Northern Europe, Pacific Islands, South-Eastern Asia, Southern America, Southern Asia, Southern Europe, Sub-Saharan Africa, and Western Europe

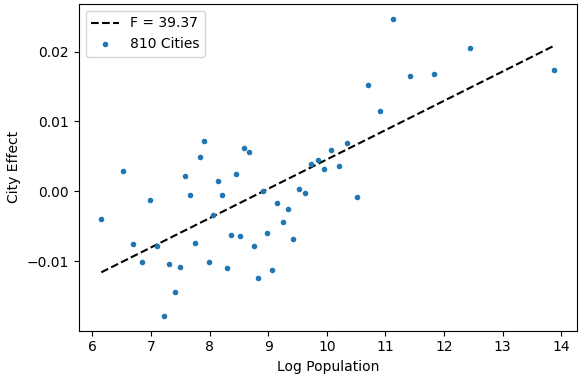
Figure 16: Correlates with City Effects



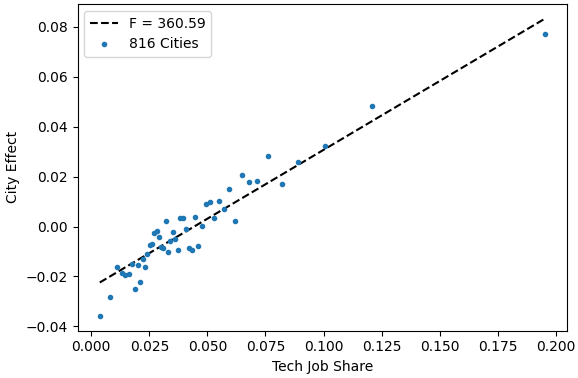
(a) Correlates of City Effects



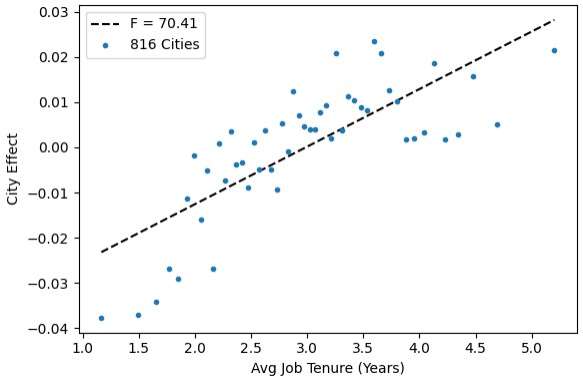
(b) Industrial Diversity



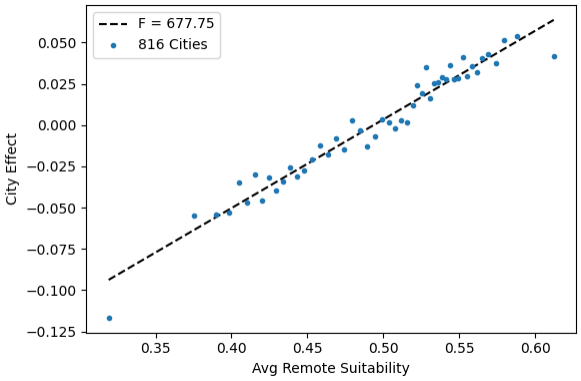
(c) Population



(d) Tech Job Share



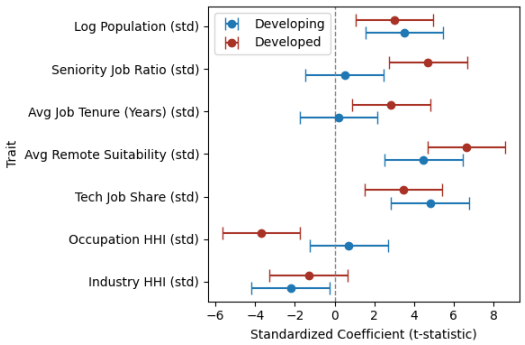
(e) Average Job Tenure



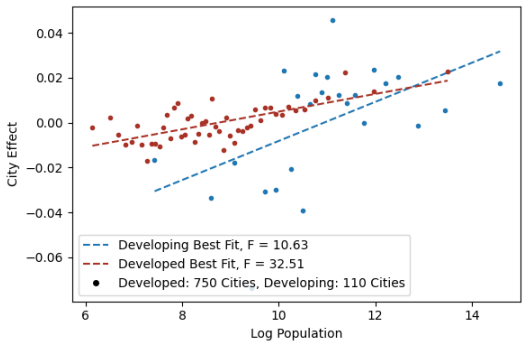
(f) Job Remote Suitability

Notes: We plot the city-level correlation between city characteristics and city effects with the inclusion of country 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.

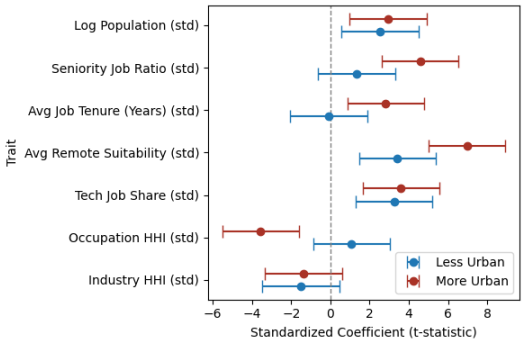
Figure 17: Correlates with City Effects By Country Characteristics



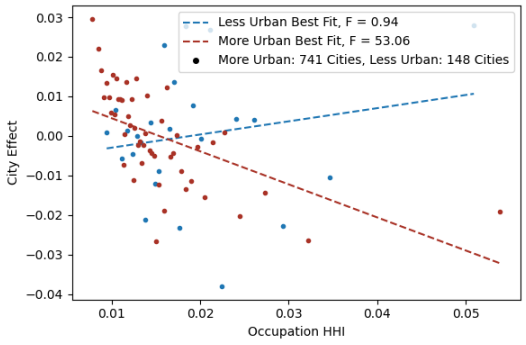
(a) By Development Status



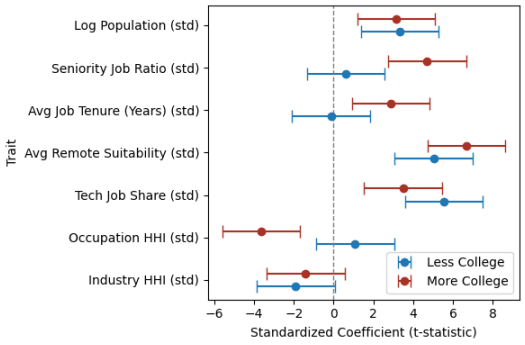
(b) Population



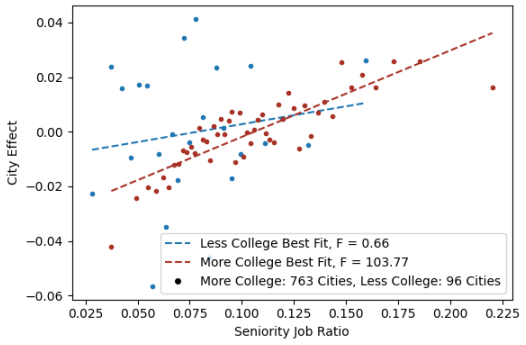
(c) By Urbanization



(d) Occupation HHI



(e) By College Attainment



(f) Seniority Ratio

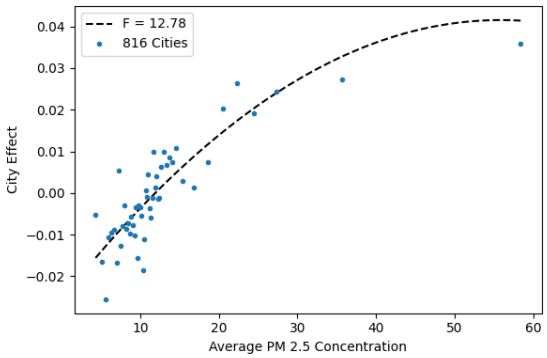
Notes: We first split our sample between countries with above/below median values for (i) GDP per capita (ii) urbanization rates and (iii) college attainment rates. We then plot the city-level correlation between city characteristics and city effects with the inclusion of country fixed effects for each 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.

set of industries, and maintain skilled, stable jobs likely see productivity gains.

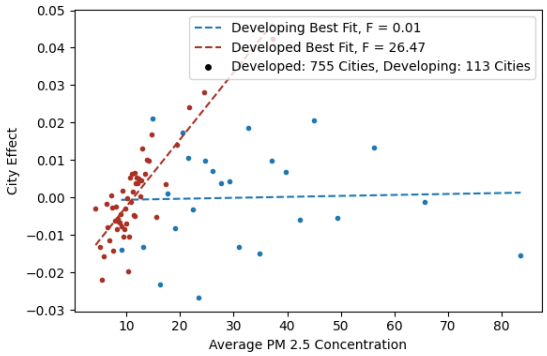
6.3 Amenities and Compensating Differentials

Our estimates are on earning gains, rather than welfare gains. In order to interpret city effects in a broader welfare context, we examine how different types of disamenities relate to productivity across cities. These disamenities fall into three conceptual categories: (a) by-products of economic production, such as pollution; (b) factors that directly affect productivity, such as pollution, precipitation or other weather shocks; and (c) factors that are largely unrelated to production and instead operate through compensating differentials, such as broader climate amenities.

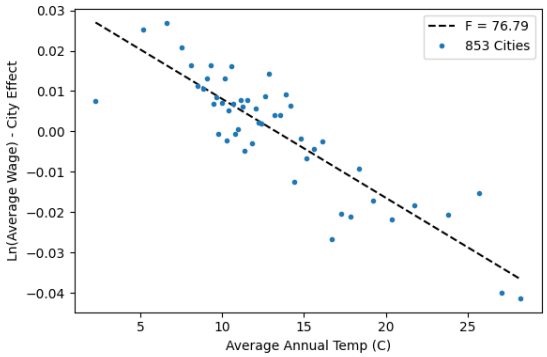
Figure 18: Pollution and City Effects



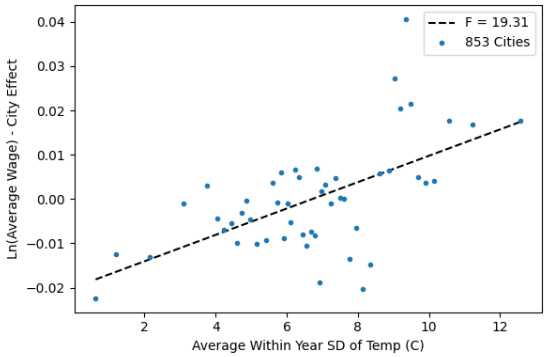
(a) Annual PM 2.5 Concentration



(b) Annual PM 2.5 Concentration, By Development



(c) Average Annual Temp



(d) Within Year SD of Temp

Notes: We plot the city-level correlation between (dis)amenities and city effects with the inclusion of country fixed effects, along with the associated F-statistics from the detailed regression. Panel (a) plots the relationship between city effects and the annual PM 2.5 concentration over the last 20 years. Panel (b) plots the relationship between city effects and the measure in (a), but with cities split by country development status. Panel (c) plots the relationship between city effects and the average annual temperature over the last 20 years. Panel (d) plots the relationship between city effects and the average within-year standard deviation of temperatures over the last 20 years.

We begin with environmental externalities that are likely generated by economic production. Pollution is a natural starting point. Figure 18a plots the relationship between annual PM 2.5 and our city effects, while Figure 18b presents the same relationship split by country development status. The most productive cities tend to have higher pollution levels—a clear disamenity and a

proxy for congestion. Notably, this relationship is driven almost entirely by cities in more developed countries. In developing countries, where pollution is uniformly high across cities, productivity is essentially uncorrelated with pollution.[18]

We then turn to environmental conditions that may influence welfare but are not necessarily tied to production. To measure compensating differentials, we examine the difference between average wages in a city and the city effect itself; locations where this measure is positive, may indicate a potentially positive compensating differential. Figures 18c and 18d plot the relationship between these potential compensating differentials and average annual temperatures, and the typical within-year variation in temperature, respectively. We find strong evidence consistent with compensating differentials globally: in cold cities, workers are typically paid more than their city effect alone predicts, while more temperate locations see the opposite pattern. Cities with more within-year temperature swings also display positive compensating differentials. However, we find no evidence of similar patterns when examining precipitation or our measure of pollution, suggesting that these latter factors either do not generate compensating differentials or are already embedded in the productivity–pollution relationship documented above.

Overall, these results highlight that disamenities associated with cities arise through multiple channels—some as by-products of production, some affecting productivity, and others shaping welfare through compensating differentials.

6.4 Spatial Decay and Agglomeration Shadows

Beyond city size and economic structure, we find that geography itself plays a significant role in explaining urban productivity. Figure 19a shows a strong negative relationship between a city's estimated effect and its distance from the nearest city within the top 10% of city effects in the same country. This pattern is consistent with modern quantitative spatial models, in which better market access enhances local productivity. Importantly, the relationship is not driven solely by proximate suburbs of major cities but also holds for smaller, independent urban centers. Although identifying "agglomeration shadows" empirically is challenging (Hornbeck et al., 2025), we find suggestive evidence that city effects begin to rise again for locations very far from the country's most productive city, consistent with reduced competition and localized agglomeration gains.

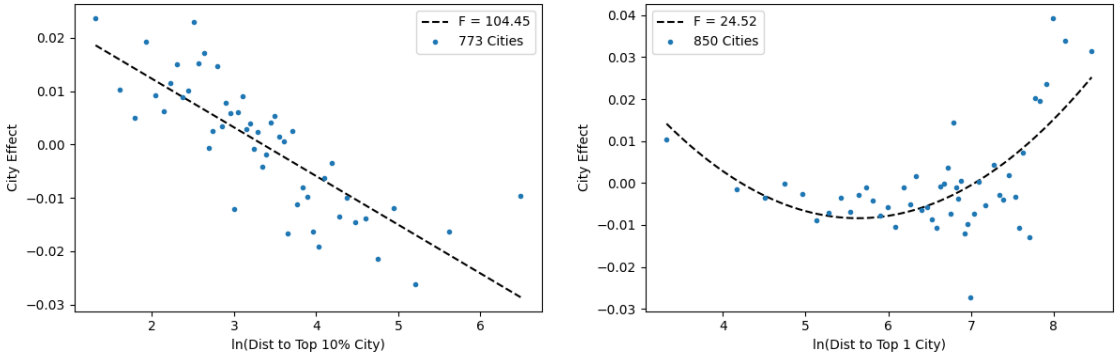
7 Aggregate Gains from Allocation Across Cities

7.1 Distance Moved, and Gains in City Effects

Our data allow us to measure the distance individuals move, both within and across countries. We test whether longer-distance moves are associated with larger productivity gains. If longer

[18] We find similar patterns when relating city effects to total production (population × wages), consistent with pollution in poorer countries being pervasive rather than concentrated in their most productive locations.

Figure 19: Distance to Big Cities and City Effects



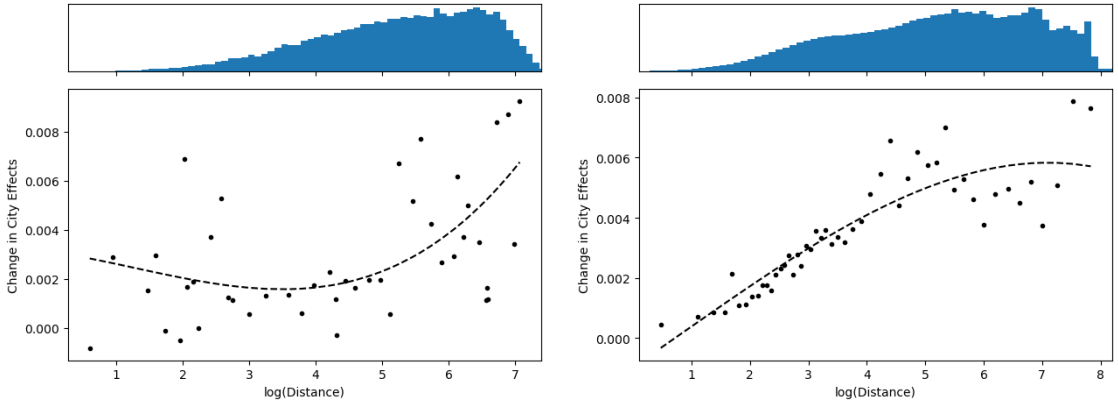
(a) Distance to Top 10% of City Effects (b) Distance to Highest City Effect

Notes: We plot the city-level correlation between city characteristics and city effects with the inclusion of country fixed effects, along with the associated F-statistics from the detailed regression. Panel (a) plots the relationship between city effects and the distance to the nearest city with a city effect in the top 10% of city effects within the country. Panel (b) plots the relationship between city effects and the distance to the city with the highest city effect in the country. Locations within 2.7km of the high-effect city are not included.

moves yield higher gains yet few individuals undertake them, this may reflect strong preferences for remaining near one's origin (a "home bias") or high migration costs that increase with distance.

The top panels of Figure 20 show the distribution of migration distances. In both low- and high-income countries, internal migration is skewed toward longer-distance moves, suggesting that fixed costs of moving are substantial—conditional on moving, individuals may choose destinations that offer the largest productivity gains. As discussed in Section 6.4, these high-productivity cities are often farther away, consistent with the presence of "agglomeration shadows."

Figure 20: Move Distance Against City Effect Gain



(a) Within Low Income Countries (b) Within High Income Countries

Distance of move (in miles) against the change in the city effect for internal moves. Atop each plot is a histogram with frequency for the move distance.

The bottom panels reveal distinct patterns across income levels. In low-income countries, moves

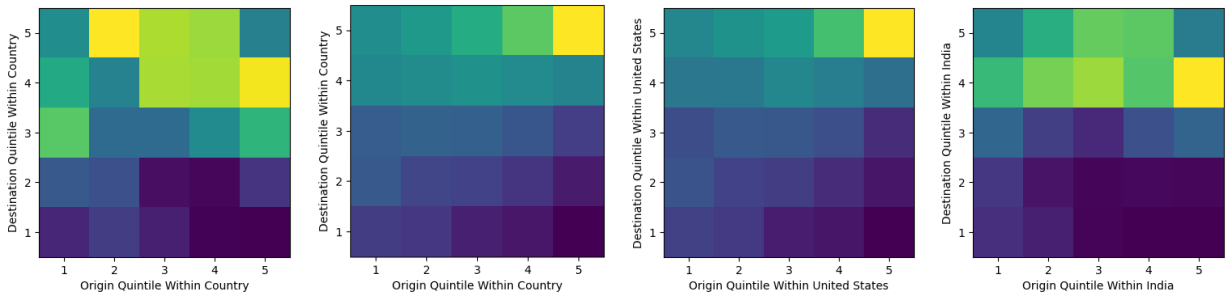
yield meaningful gains even over short distances, with much larger gains at very long distances. In high-income countries, by contrast, the relationship between distance and city-effect gains is strongly upward sloping—each additional mile moved corresponds to higher potential productivity gains. These patterns suggest that while individuals in all contexts are more likely to move long distances and thereby capture larger returns, geographic preferences and regional attachments (e.g., remaining in one's home region) may prevent many from realizing the most productive opportunities available within their country.

7.2 Transitions Across the City Effect Distribution

There are substantial potential gains from moving across cities. Such mobility can improve the allocation of labor and raise aggregate earnings when individuals relocate from low- to high-productivity cities. To assess this mechanism, we revisit our transition matrices (previously constructed across *wage* quintiles) by now analyzing transitions across the *city-effect* distribution.

The top panel of Figure 4 above already presented transition matrices for internal migration by city wage quintiles. Across most countries, the brightest cells appeared in the top row, indicating that individuals were disproportionately moving to the highest-wage cities, regardless of where they started in the wage distribution. This pattern held not only in high-income countries such as the United States but also in lower-income settings like India. If higher wages accurately reflect greater productivity, these moves would generate meaningful allocative gains.

Figure 21: Transition Matrices for Internal Movers by City Effect Quintile



(a) Low Income Countries (b) High Income Countries (c) In the US (d) In India

Notes: The figures show the transition matrices for within-country moves between cities by city effect quintiles. Restricted only to movers between cities with 500 or more users in the regression sample.

Figure 21, in contrast, is based on city effect quintiles: the y-axis denotes the destination city-effect quintile, and the x-axis the origin. It shows that wages do not always align with productivity, implying unrealized gains from reallocation. This disconnect is especially pronounced in low-income countries (Figure 21a), where many migrants move to cities in the third and fourth quintiles of the city-effect distribution rather than to the most productive locations. In high-income countries, where wages and city effects are more closely correlated, most moves still target cities with the

highest city effects—though a nontrivial share of moves also go to fourth-quintile cities.

In India, most internal moves are to cities in the fourth quintile of the city-effect distribution. Further investigation reveals that while these tend to be large, high-wage urban centers, the most productive locations are their surrounding satellite towns. These contrasting patterns highlight that in low-income countries, particularly those with uneven spatial development, there remain sizable potential earnings gains from redirecting migration to the truly high-productivity locations.

7.3 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 within a country were uniform across countries, individuals moving from a city in the 25th to 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).

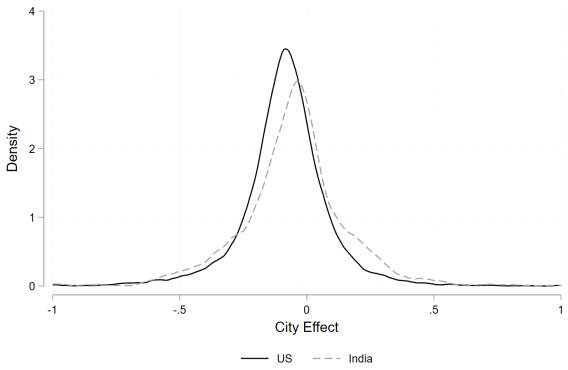
In order to highlight the differences in the distribution of city effects across countries, Figures 22a-22c plot the city effect distributions in the US and India, along with a sample of developed and developing countries. Two important patterns immediately emerge: first, the distribution of city effects 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.

A natural next step is determining migration's role in shaping the share of wage differences explained by city effects. If migration barriers are high, individuals are less able to efficiently sort across space. The importance of city effects would thus be higher in such an environment, as 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 variation in city effects within a country.[19] Figure 22d reveals a strong negative relationship between the internal migration rate of a country and the variance of city effects.

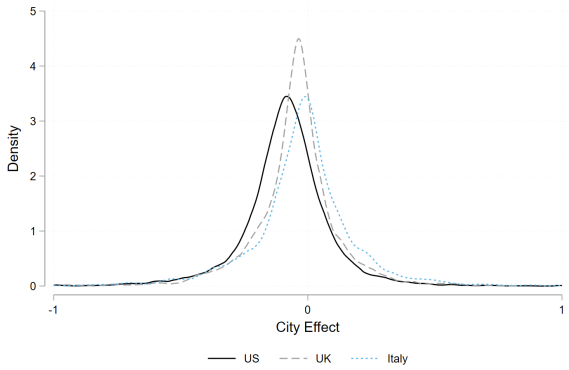
Indeed, our findings indicate that the potential returns to internal migration differ considerably between countries, depending on the underlying distribution of city effects. For instance, countries with a higher dispersion of city effects–often developing countries–may offer greater economic incentives 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.

[19] Bell et al. (2015) covers 96 countries. Practically, we limit our sample of countries to those with a coverage rate above 0.1, more than 50 identified cities, and over 20,000 unique users that appear within that country.

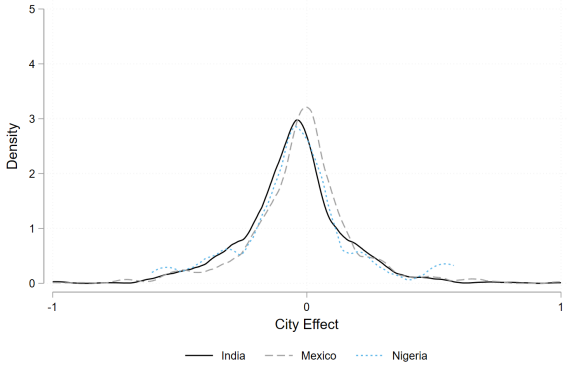
Figure 22: Distribution of City Effects



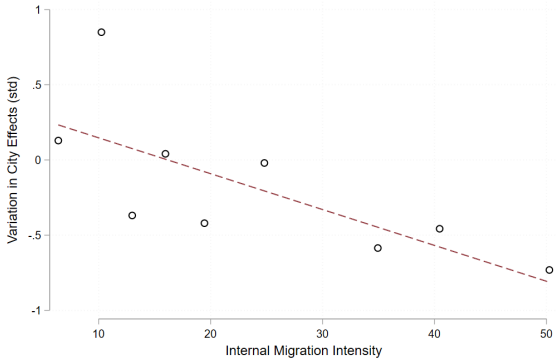
(a) City Effects Distribution, USA and India



(b) City Effects Distribution, Developed Countries



(c) City Effects Distribution, LMICs



(d) Migration Intensity and Variation in City Effects

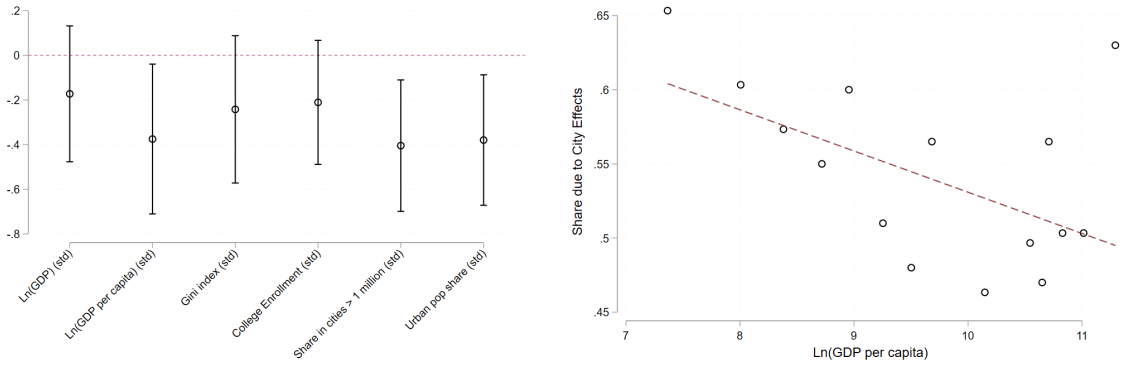
Panels (a)-(c) plots the density plot of the standard deviation of city effects estimated using Equation 3 across countries. The distribution indicates that the variance of city effects varies meaningfully across countries. Panel (a) shows the distribution of city effects for the USA and India. We recenter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

7.4 Explaining the Importance of City Effects by Country

Finally, we examine what explains the importance of city effects across 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 23a plots the estimated coefficients for the relationship between the share of wage differences explained by city effects with each of our explanatory variables while controlling for the coverage rate for each country.

Overall, we find that wealthier, educated, and urbanized countries exhibit a smaller share of wage variation explained by city effects. As a higher share attributed to city effects implies greater potential gains from reallocating workers from low- to high-productivity cities, this pattern suggests

Figure 23: Determinants of the Importance of City Effects



(a) Country Aggregates and Median Share

(b) GDP Per Capita and Median Share

Panel (a) plots the coefficients along with 95% CI from regressions on a variety of variables on the share of wage differences between above and below-median wage cities that is due to differences in city effects within a country. Each variable is standardized, and robust standard errors are implemented. Panel (b) plots the relationship between the median share in (a) and country level GDP per capita.

the returns to internal migration are likely larger in low- and middle-income countries. This finding is consistent with the idea that spatial frictions are more pronounced in developing economies.

Figure 23b illustrates a strong negative relationship between GDP per capita and the share of wage differences explained by city effects—particularly among low- and middle-income countries. But the relationship weakens among high-income economies, possibly reflecting the growing role of housing supply constraints and land-use regulations (Hsieh and Moretti, 2019). While the largest potential income gains from reducing spatial barriers are in poorer countries, our results also indicate that alleviating mobility frictions could yield meaningful gains even in advanced economies.

8 Quantifying Gains from Reallocation Within and Across Cities

Finally, we turn to quantifying the potential wage gains from reallocating workers based on our results. There are two natural exercises that our results allow us to examine. The first is to estimate the increase in aggregate wages that would result from reallocating workers across cities within each country to match the observed relationship between city effects and city size in the US.

We first rank cities by city effects within a country and then compute the resulting CDF that captures how the population is distributed across city rank.[20] We then reallocate each country's population to match the US CDF. Figure 24a compares the observed and counterfactual distributions for India against the US. Several large Indian cities fall just above the 65th percentile of the distribution rather than near the top, indicating productivity potential from reallocation.

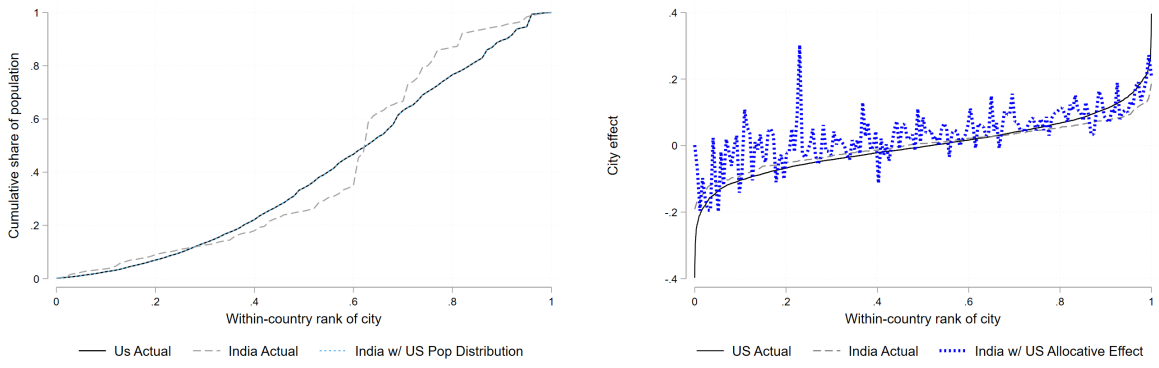
A key implication of reallocating individuals across cities is that the average estimated individual

[20] Given cross-country differences in the number of cities and the variation in estimated city effects, we group cities into 100 rank-based bins for the counterfactual analysis.

productivity changes as workers from less productive cities move to locations where the typical worker is more productive. We assume that, prior to reallocation, all individuals within a city share the city's average individual productivity. This assumption allows us to randomly assign migrants from origin to destination cities and recalculate expected average productivity after reallocation. Importantly, this exercise holds both the national average of individual productivity and the city effects fixed, so any wage gains arise solely from relocating individuals to more productive cities. Even under this conservative setup, the potential gains are substantial: average annual wages in India rise by 2.3% (about USD 200 per worker) following reallocation, while in other developing countries such as Mexico and Nigeria, the potential gains reach 4.0% and 5.9%, respectively.

Our second quantification exercise draws on our decomposition results in Section 6.1, where we show that in countries such as the US, workers are better allocated to the most productive firms *within* a location. Rather than calculating the potential gains from moving individuals across cities, we now assess the potential gains from simply moving the same number of workers across firms within cities. We group cities in the US based on their city effects in the same way as in the previous exercise. Then, we calculate the average value of the allocative effect for city groups in the US and apply those values to the mean effect, $\bar{\gamma}_j$ for cities in other countries.[21]

Figure 24: Quantifying the Gains from Re-allocation



(a) Population Reallocation

(b) Within-City Reallocation

Panel (a) plots the actual CDF of population by city effect rank for the US and India, along with the counterfactual CDF for India when we carry out the population redistribution exercise to match the US distribution Panel (b) plots the relationship between city effects and the within-country rank of city effects for the actual US and Indian city effects, as well as the counterfactual Indian city effects once we add the corresponding allocative effect from the US to the average firm effect within Indian cities.

Figure 24b plots the observed and counterfactual distribution of city effects for India against the US; the fact that counterfactual city effects are almost uniformly higher than observed values in India is evidence that the average allocative effect in US cities is higher than in India. Additionally, we see relatively large increases in city effects for India at the bottom and top of the city effect

[21] In practice, this adjustment yields smaller city effects than the mean effects for less productive cities and larger ones for the most productive cities. Whether these counterfactual city effects exceed or fall below our estimated values depends on the relative magnitude of allocative effects across countries.

distribution, highlighting the potential for within-city reallocation to improve allocation within a country. We find slightly larger potential gains from this within-city reallocation exercise, with India increasing average wages by 2.6%. However, the gains are not uniform across countries. In Mexico and Nigeria, the potential wage increases from within-city reallocation are smaller—1.6% and 2.9%, respectively—than those from across-city reallocation. Consistent with expectations about the diminishing scope for firm-level allocative improvements over the development process, the gains are modest in high-income countries such as the UK (0.9%) and Germany (1.0%).

Our final exercise combines the two counterfactuals above. We first modify city effects based on the allocative effect of the US as described above, and then reallocate population according to the modified city effects to match the US distribution. Although these two mechanisms could offset one another—if improvements in one reduce the gains from the other—we find little evidence of substitution. For India, the combined effect raises average wages by 4.3%, with roughly 90% of the gains explained by the sum of the two exercises considered separately. The corresponding figures are slightly higher for Mexico and Nigeria. These results suggest that development policies aimed at reducing spatial and firm-level frictions are not competing, and that addressing both dimensions simultaneously can yield substantial aggregate income gains.

9 Conclusion

We provide 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. Our event studies isolate the impact of city effects on individual wages, job seniority, and occupational and industrial scores. We disentangle 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 low-income countries have more low-productivity cities, particularly because a higher share of workers work in less productive firms. Our city effects are associated with population, industrial diversity, and the prevalence of high-skill jobs, reinforcing the importance of agglomeration economies. Notably, the variance in city effects decreases with a country's 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 productivity interacts with national and regional contexts to shape labor markets. 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 city productivity.

References

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

Abramitzky, Ran, Leah Platt Boustan, and Katherine Eriksson (2016) "Europe's Tired, Poor, Huddled Masses: Self-Selection and Economic Outcomes in the Age of Mass Migration," *American Economic Review*, 102 (5), 1832–1856.

Abramitzky, Ran, Leah Platt Boustan, Santiago Perez, and Elisa Jacome (2021) "Intergenerational Mobility of Immigrants in the US over Two Centuries," *American Economic Review*, 111 (2), 580–608.

Akbar, Prottoy, Victor Couture, Gilles Duranton, and Adam Storeygard (2024) "The Fast, the Slow, and the Congested: Urban Transportation in Rich and Poor Countries," *Working Paper*.

Albert, Christoph and Joan Monras (2022) "Immigration and Spatial Equilibrium: The Role of Expenditures in the Country of Origin," *American Economic Review*, 112 (11), 3763–3802.

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

Asher, Sam, Kritarth Jha, Paul Novosad, Anjali Adukia, and Brandon Tan (2025) "Residential Segregation and Unequal Access to Local Public Services in India: Evidence from 1.5 Million Neighborhoods,"Technical report, Working Paper.

Badinski, Ivan, Amy Finkelstein, Matthew Gentzkow, and Peter Hull (2023) "Geographic Variation in Healthcare Utilization: The Role of Physicians," *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.

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.

Bollinger, Christopher R. and Barry T. Hirsch (2006) "Match bias from earnings imputation in the Current Population Survey: The case of imperfect matching," *Journal of Labor Economics*, 24 (3), 491–528, 10.1086/504276.

——— (2013) "Is earnings nonresponse ignorable?" *The Review of Economics and Statistics*, 95 (2), 407–416, 10.1162/REST_a_00252.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2019) "A Distributional Framework for Matched Employer–Employee Data," *Econometrica*, 87 (3), 699–739, 10.3982/ECTA15722.

Borovičková, Katarína and Robert J. Shimer (2024) "Assortative Matching and Wages: The Role of Selection,"Technical report, National Bureau of Economic Research, Working Paper No. 33184.

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

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 (2025) "Location, location, location," *American Economic Journal: Applied Economics*, 17 (1), 297–336.

Carry, Pauline, Benny Kleinman, and Elio Nimier-David (2025) "Location Effects or Sorting? Evidence from Firm Relocation," *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.

Clemens, Michael A. (2011) "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?" *Journal of Economic Perspectives*, 25 (3), 83–106, 10.1257/jep.25.3.83.

Clemens, Michael A., Claudio E. Montenegro, and Lant Pritchett (2019) "The Place Premium:

Bounding the Price Equivalent of Migration Barriers," *Review of Economics and Statistics*, 101 (2), 201–213, 10.1162/rest_a_00718.

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.

Couture, Victor and Jessie Handbury (2020) "Urban Revival in America," *Journal of Urban Economics*, 119, 103267, 10.1016/j.jue.2020.103267.

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

Desmet, Klaus and Esteban Rossi-Hansberg (2013) "Urban Accounting and Welfare," *American Economic Review*, 103 (6), 2296–2327, 10.1257/aer.103.6.2296.

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.

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 and Philipp Kircher (2011) "Identifying Sorting—In Theory," *The Review of Economic Studies*, 78 (3), 872–906.

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

Eli, Shari, Laura Salisbury, and Allison Shertzer (2016) "Migration Responses to Conflict: Evidence From the Border of the American Civil War," *Working Paper*.

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.

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.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010) "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings," *Journal of Political Economy*, 118 (3), 536–598.

Hanna, Rema, Gabriel Kreindler, and Benjamin A. Olken (2017) "Citywide Effects of High-Occupancy Vehicle Restrictions: Evidence from "Three-in-One" in Jakarta," *Science*, 357 (6346), 89–93, 10.1126/science.aal2036.

Harari, Mariaflavia (2020) "Cities in Bad Shape: Urban Geometry in India," *American Economic Review*, 110 (8), 2377–2421, 10.1257/aer.20181459.

——— (2024) "Residential Patterns and Local Public Goods in Urban Brazil,"Technical report, Working Paper.

Harari, Mariaflavia and Maisy Wong (2025a) "Slum Upgrading and Long-run Urban Development: Evidence from Indonesia," *Review of Economic Studies*, Forthcoming.

——— (2025b) "Colonial Legacy and Land Market Formality," *Journal of Urban Economics*, 149, 103789, 10.1016/j.jue.2025.103789.

Heise, Sebastian and Tommaso Porzio (2021) "The Aggregate and Distributional Effects of Spatial

Frictions," NBER Working Paper 28792, National Bureau of Economic Research.
Hornbeck, Richard and Daniel Keniston (2017) "Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872," *American Economic Review*, 107 (6), 1365–1398.
Hornbeck, Richard, Guy Michaels, and Ferdinand Rauch (2025) "Identifying Agglomeration Shadows: Long-run Evidence from Ancient Ports," *Working Paper*.
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.
Hsieh, Chang-Tai and Enrico Moretti (2019) "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 11 (2), 1–39.
Khanna, Gaurav, Wenquan Liang, Ahmed Mushfiq Mobarak, and Ran Song (2025) "The Productivity Consequences of Pollution-Induced Migration in China," *American Economic Journal: Applied Economics*, 17 (2), 184–224, 10.1257/app.20220655.
Kline, Patrick M. (2024) "Firm Wage Effects,"Technical report, National Bureau of Economic Research Working Paper No. 33084, 10.3386/w33084, Also available as Handbook of Labor Economics, Vol. 5 (2024), pp. 115–181.
Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten (2020) "Leave-out estimation of variance components," *Econometrica*, 88 (5), 1859–1898.
de La Roca, Jorge and Diego Puga (2017) "Learning by working in big cities," *The Review of Economic Studies*, 84 (1), 106–142.
Martellini, Paolo, Todd Schoellman, and Jason Sockin (2024) "The Global Distribution of College Graduate Quality," *Journal of Political Economy*, 132 (2), 434–483.
McKenzie, David, John Gibson, and Steven Stillman (2010) "How Important is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration," *Journal of the European Economic Association*, 8 (4), 913–945, 10.1111/j.1542-4774.2010.tb00544.x.
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.
Olley, G. Steven and Ariel Pakes (1996) "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64 (6), 1263–1297.
Orefice, Gianluca and Giovanni Peri (2025) "Immigrants and the Efficiency of Labor Allocation," *Review of Economics and Statistics*, Forthcoming.
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.
The Occupational Information Network (2020) "The Occupational Information Network (O*NET) Database 14.0," https://www.onetcenter.org/db_releases.html, Accessed May 26, 2025.
U.S. Census Bureau (2022) "North American Industry Classification System (NAICS)," https://www.census.gov/naics/, Accessed May 26, 2025.
Young, Alwyn (2013) "Inequality, the Urban-Rural Gap, and Migration," *The Quarterly Journal of Economics*, 128 (4), 1727–1785, 10.1093/qje/qjt025.

A Appendix

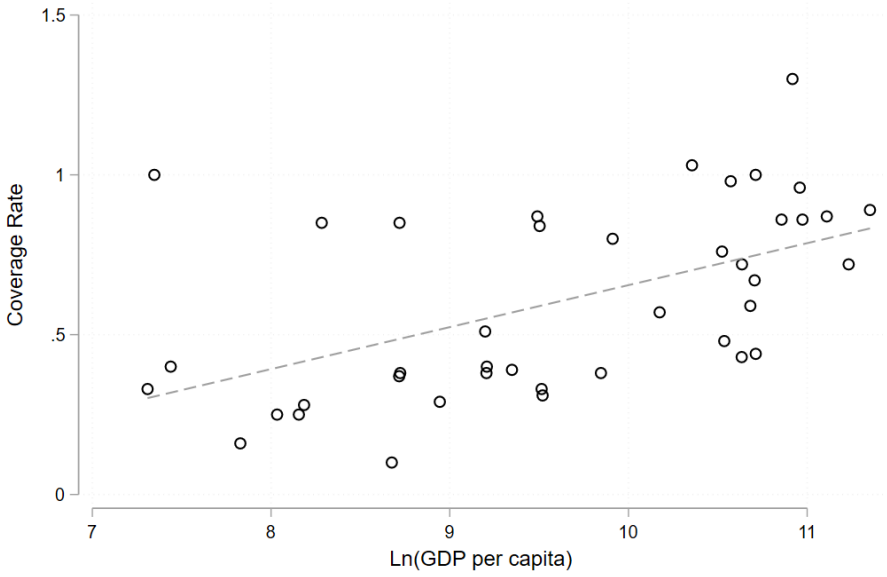
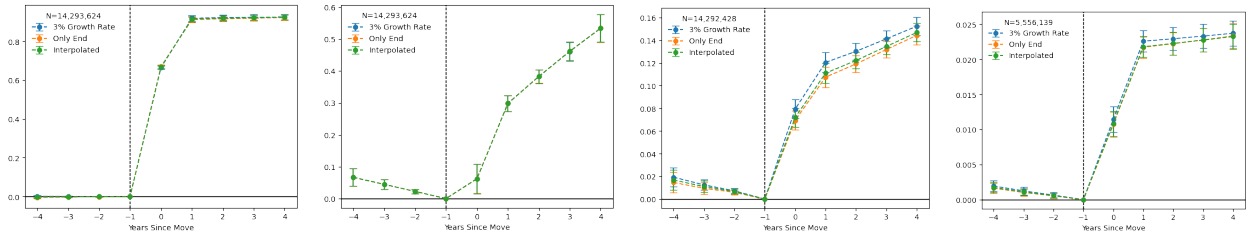


Figure A1: Coverage Rate, by GDP per Capita

Notes: Using our final sample to make cross-country comparisons, we plot each country's coverage rate as defined by Amanzadeh et al. (2024) against ln(GDP per capita) along with a line of best fit.

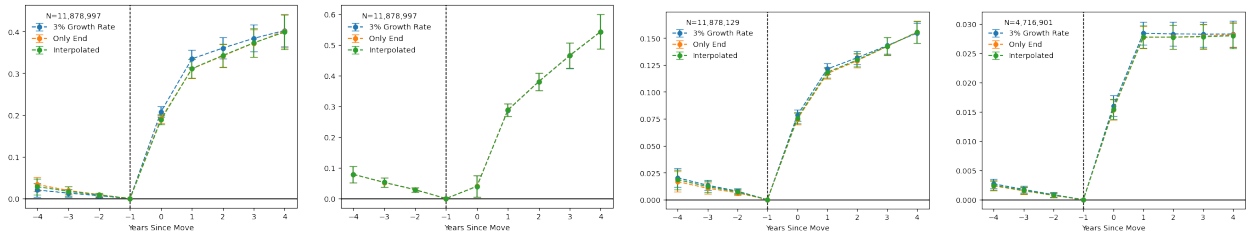
Figure A2: Event Studies by Outcome for Different Panels: Within-Country Movers and All Movers

Panel A: All Movers



(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: Within-Country Movers

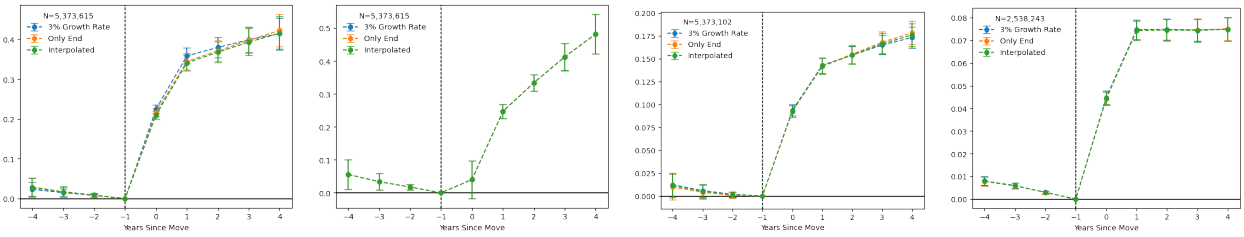


(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

Analogous to Figure 5 with differently constructed datasets. The "3% Growth Rate" panel, which is used in the paper, applies a 3% growth rate to user salaries for as long as they remain in the same position. The "Only End" panel uses only the end salary for each year a user is at a position. The "Interpolated" panel uses a log-salary interpolated between the start and end of the users stint at a position. 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.

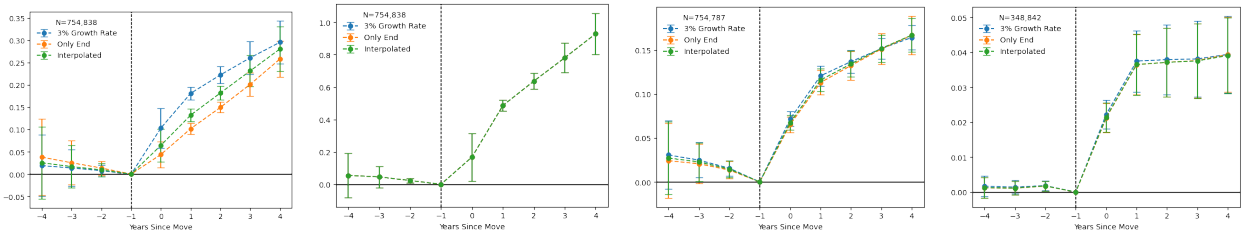
Figure A3: Event Study for Internal Movers for Different Panels, Developed Economies

Panel A: United States



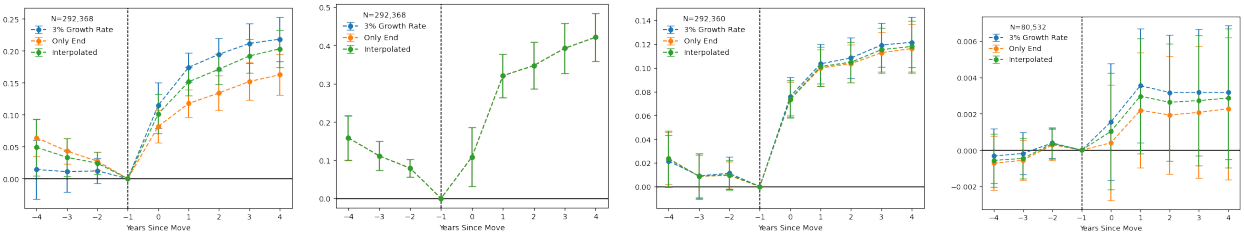
(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: United Kingdom



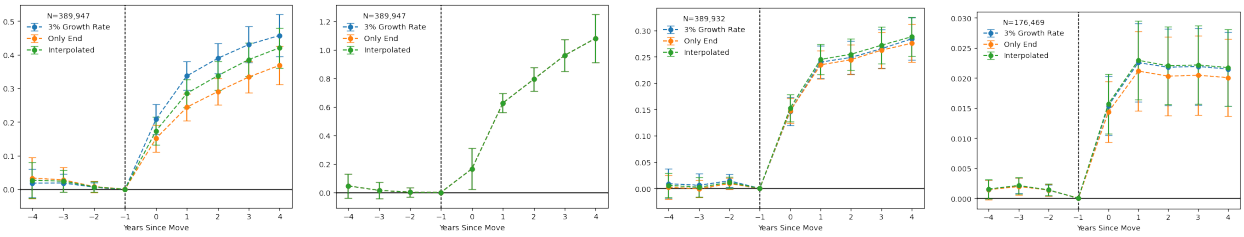
(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

Panel C: Italy



(i) Log Salary (j) Seniority (k) ONET Wage (l) NAICS Wage

Panel D: Germany

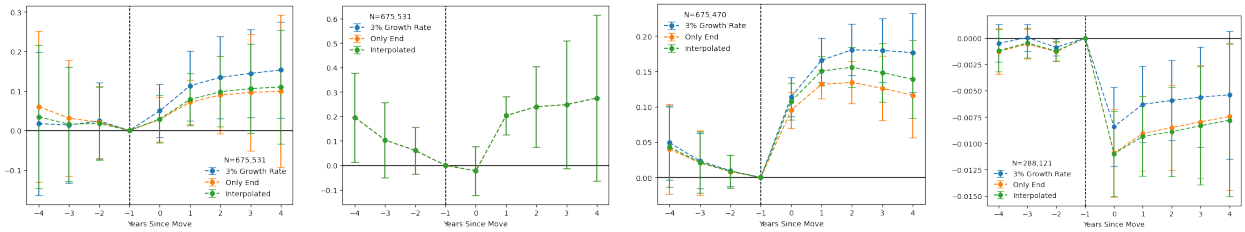


(m) Log Salary (n) Seniority (o) ONET Wage (p) NAICS Wage

Analogous to Figure 6 with differently constructed datasets. The "3% Growth Rate" panel, which is used in the paper, applies a 3% growth rate to user salaries for as long as they remain in the same position. The "Only End" panel uses only the end salary for each year a user is at a position. The "Interpolated" panel uses a log-salary interpolated between the start and end of the users stint at a position. 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.

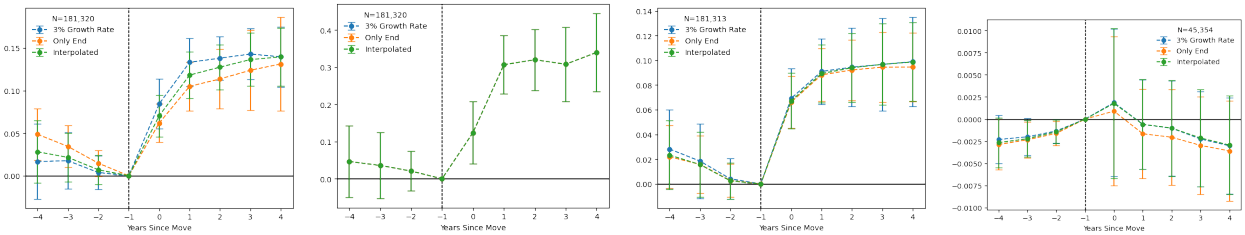
Figure A4: Event Study for Internal Movers for Different Panels, Developing Economies

Panel A: India



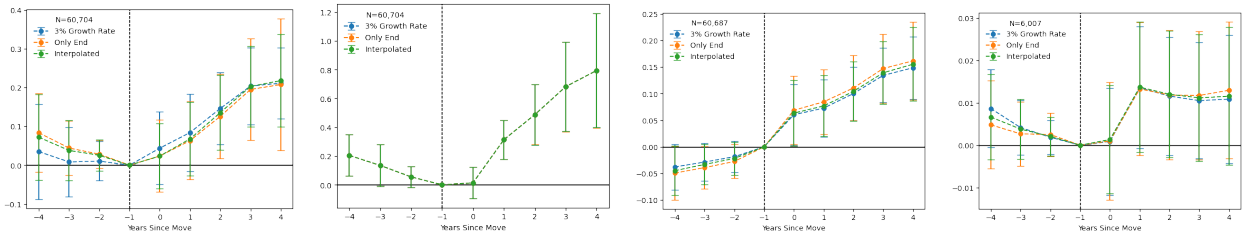
(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: Mexico



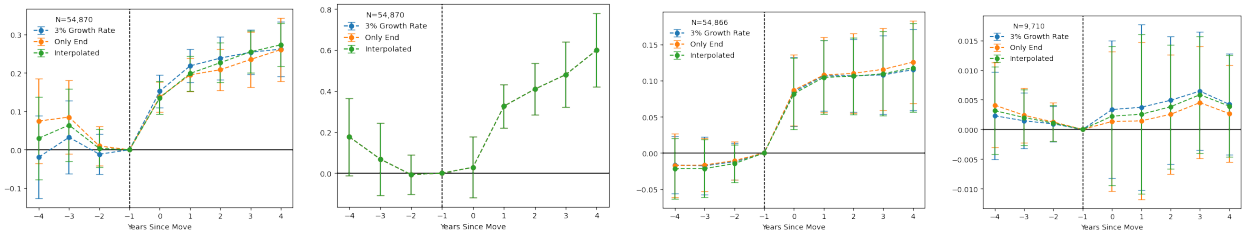
(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

Panel C: Nigeria



(i) Log Salary (j) Seniority (k) ONET Wage (l) NAICS Wage

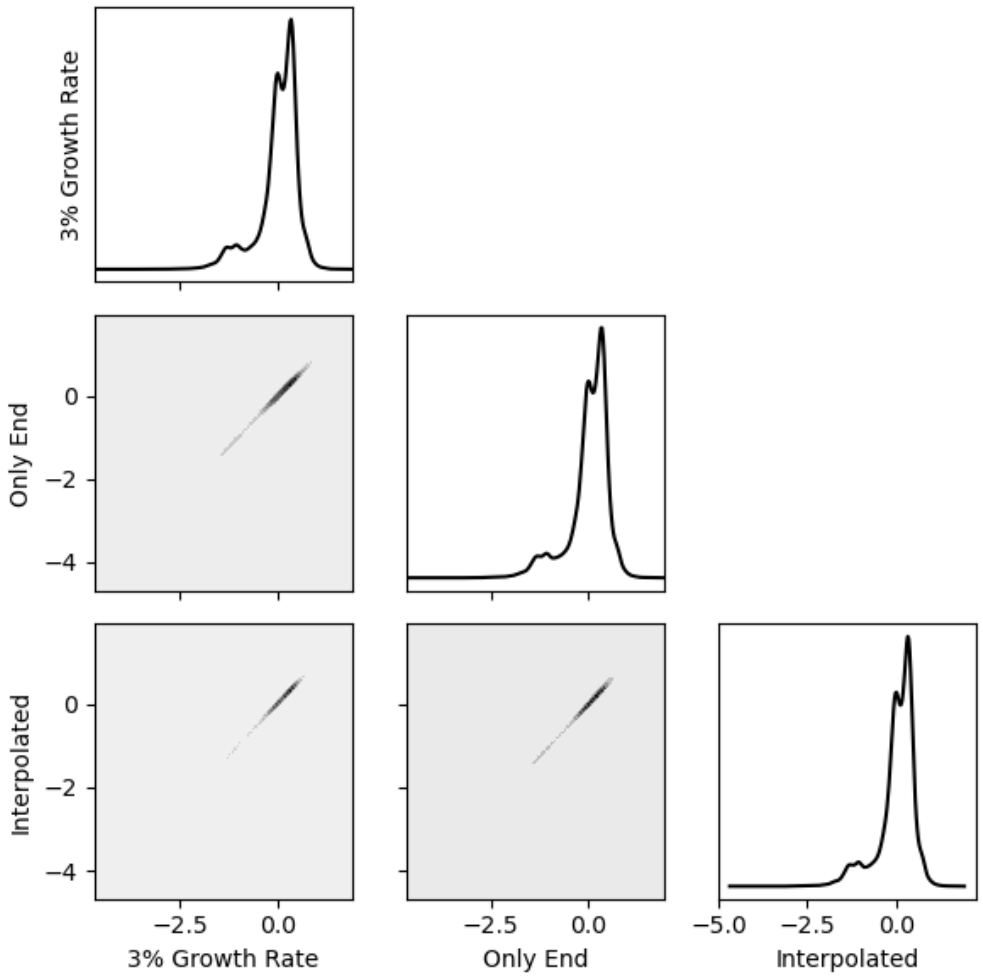
Panel D: Chile



(m) Log Salary (n) Seniority (o) ONET Wage (p) NAICS Wage

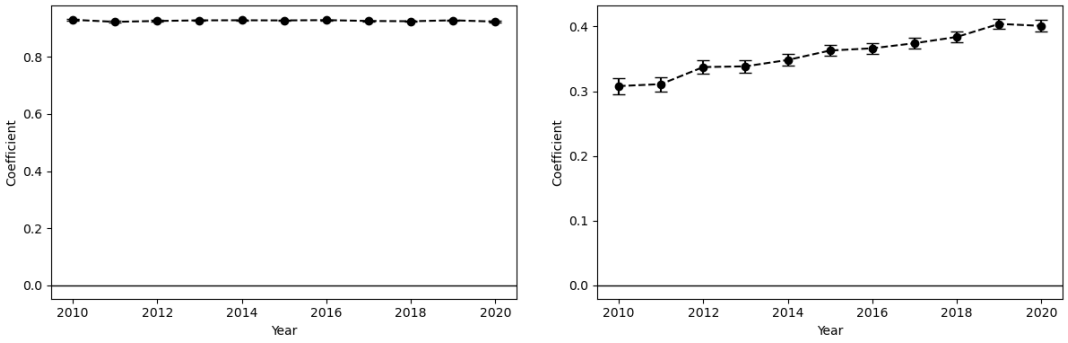
Analogous to Figure 7 with differently constructed datasets. The "3% Growth Rate" panel, which is used in the paper, applies a 3% growth rate to user salaries for as long as they remain in the same position. The "Only End" panel uses only the end salary for each year a user is at a position. The "Interpolated" panel uses a log-salary interpolated between the start and end of the users stint at a position. Both estimation samples are subset to only include individuals who worked in the year prior to moving ($t = -1$). Standard errors clustered at the user level.

Figure A5: City Effects Distribution for Different Panels



Distribution of city effects for different panels. The diagonal shows the univariate distribution of city effects. Below the diagonal shows the bivariate distribution of city effects with different panels. The "3% Growth Rate" panel, which is used in the paper, applies a 3% growth rate to user salaries for as long as they remain in the same position. The "Only End" panel uses only the end salary for each year a user is at a position. The "Interpolated" panel uses a log-salary interpolated between the start and end of the users stint at a position.

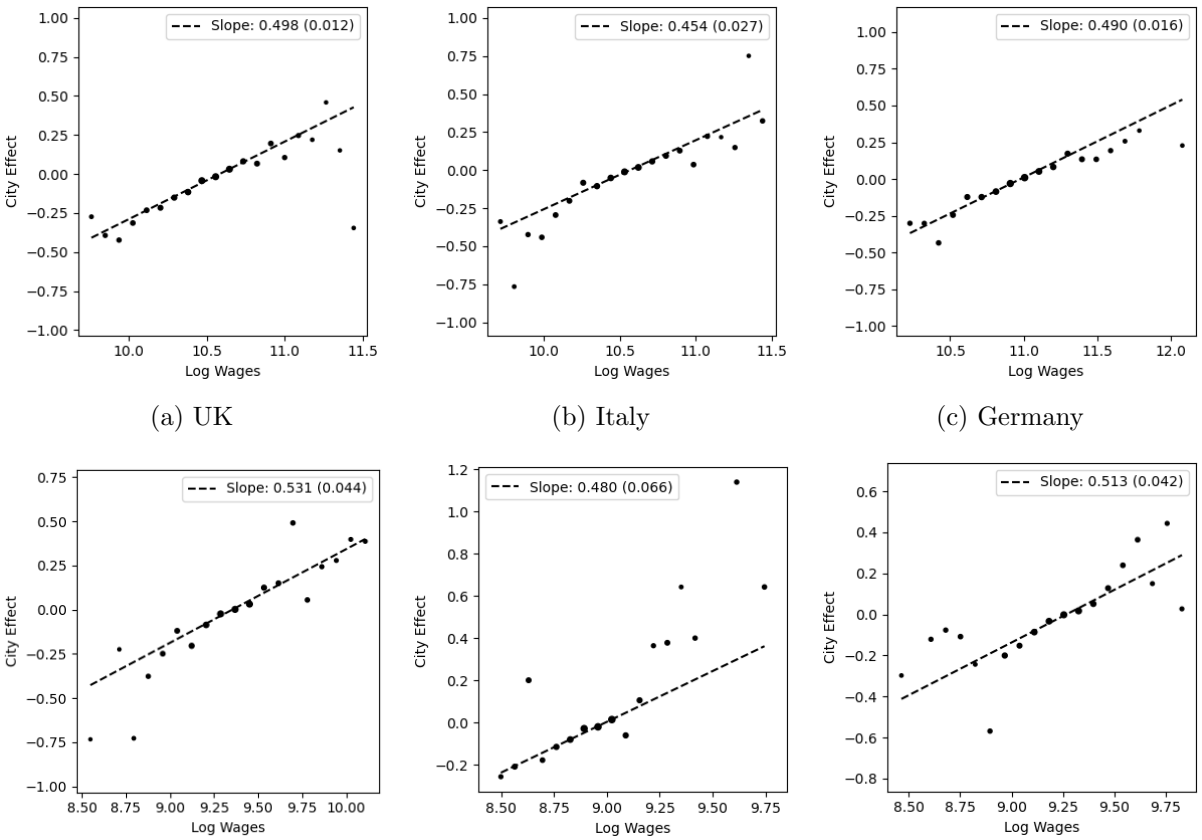
Figure A6: Changes in Wages by Difference in City Wages, Over Time



(a) International Moves (b) 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), by year of move. The left panel is for international moves, and the right panel is for within-country moves.

Figure A7: City Effects, Other Countries

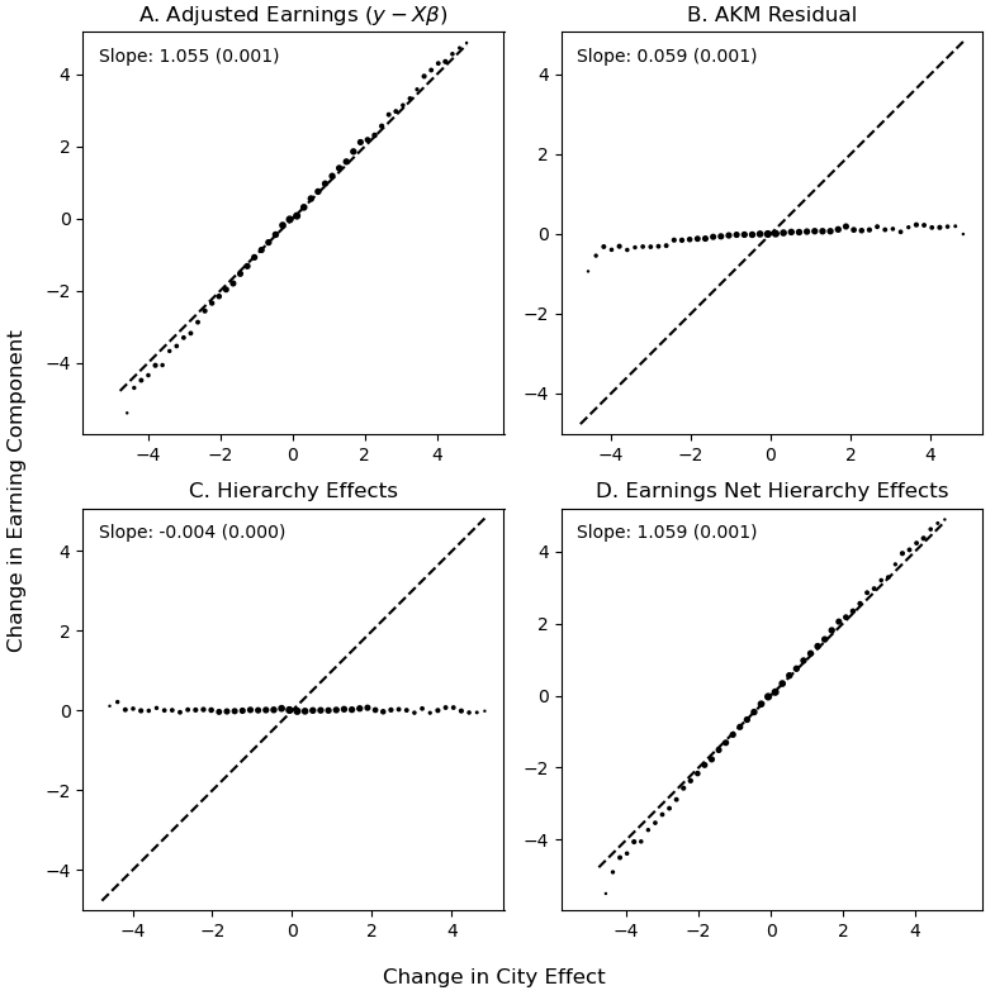


(a) UK (b) Italy (c) Germany

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

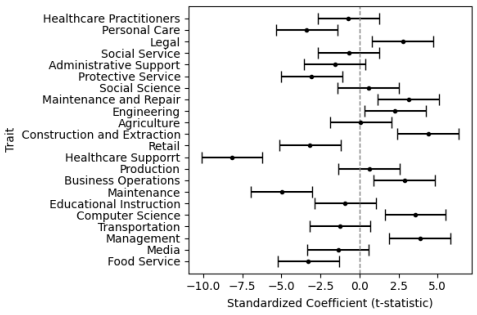
Each panel plots the relationship between the City Effect and the Log Average City Salary, weighting cities by the number of users, where points are equally spaced across the distribution and the size of each point is proportional to the fraction of users in cities in each wage bin.

Figure A8: Change in earnings components of city movers around move, against the change in city effect (all movers)

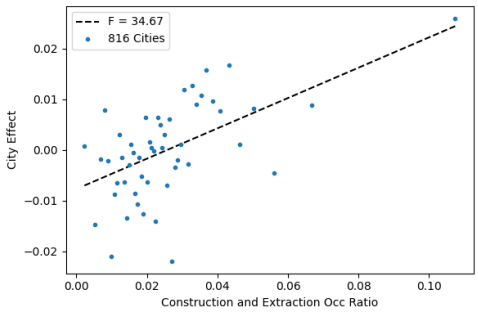


Notes: We look at changes in earnings two years before and after a move. X-axis plots the change in city effects for movers as destination city effect - origin city effect. City effects are constructed using equation 4. We plot a 45' line for reference. The sample is restricted to movers who move once.

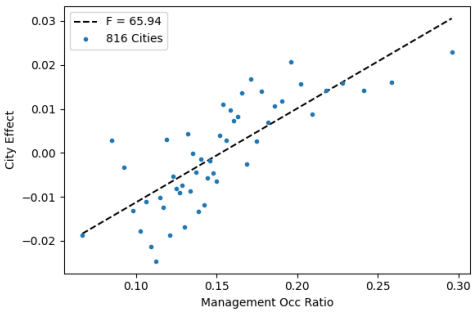
Figure A9: Correlates with City Effects By Country Characteristics



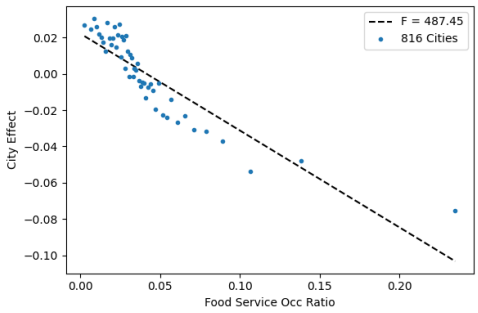
(a) Correlates of City Effects



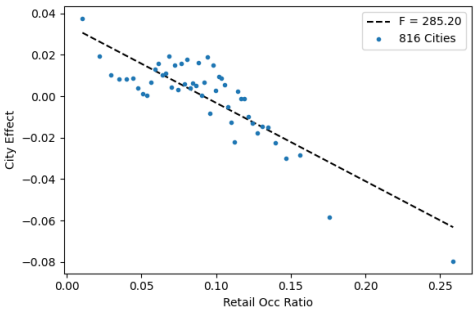
(b) Construction



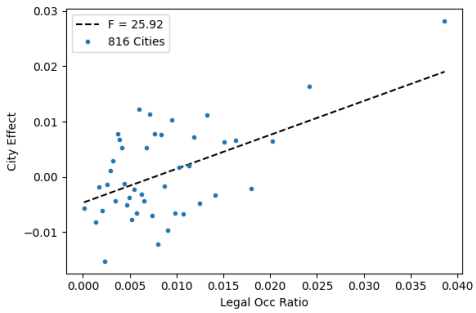
(c) Management



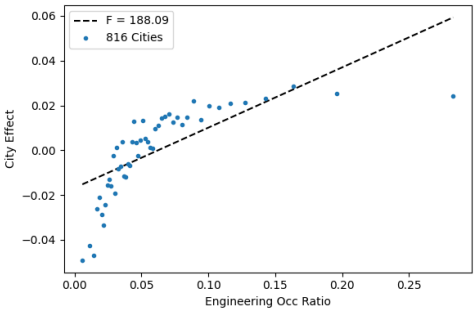
(d) Food Service



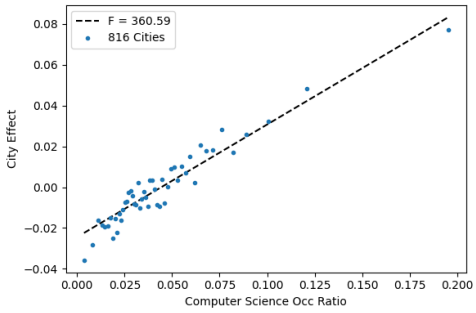
(e) Retail



(f) Legal



(g) Engineering



(h) Computer Science

Notes: We plot the city-level correlation between city characteristics (standardized two-digit SOC shares within a city) and city effects with the inclusion of country fixed effects. All correlates are two-digit SOC employment shares within a city. Standard errors are clustered at the country level.

Table A1: Dataset Summary Statistics

	Moved Firms	**Moved Cities**	**Moved Countries**	**Not in Dataset**
Panel A: All				
Salary	43,415.65	51,043.42	47,830.56	48,420.13
	(33,002.43)	(32,289.17)	(32,677.85)	(35,987.82)
NAICS Wage	51,848.30	52,026.82	52,275.66	47,883.93
	(9,604.69)	(9,058.45)	(8,803.39)	(12,982.51)
ONET Wage	48,589.38	48,074.86	50,897.55	49,356.41
	(13,172.00)	(12,567.75)	(13,431.65)	(15,439.86)
Seniority	2.48	2.39	2.63	2.67
	(1.41)	(1.41)	(1.45)	(1.61)
Num. Users	5,241,672	3,638,898	783,360	193,900,336
Panel B: United States				
Salary	63,753.10	64,052.72		63,831.08
	(36,280.31)	(32,930.04)		(39,277.46)
NAICS Wage	64,218.43	65,085.95		59,351.19
	(16,105.24)	(15,536.17)		(19,103.08)
ONET Wage	63,290.77	64,821.54		64,154.50
	(23,966.32)	(23,522.72)		(26,599.89)
Seniority	2.43	2.38		2.64
	(1.50)	(1.47)		(1.65)
Num. Users	1,644,316	1,708,471		66,721,565
Panel C: India				
Salary	10,482.96	10,518.35		10,722.45
	(3,254.26)	(3,445.25)		(4,939.91)
NAICS Wage	10,641.67	10,757.72		10,238.83
	(1,508.29)	(1,420.27)		(2,050.60)
ONET Wage	10,482.93	10,636.35		10,806.14
	(1,872.67)	(1,981.41)		(2,632.39)
Seniority	2.62	2.58		2.77
	(1.18)	(1.23)		(1.45)
Num. Users	751,519	306,936		11,958,280

Notes: Summary statistics by groups of individuals. "Moved Firms" indicate individuals who have moved firms within a city (i.e., not moved cities). "Moved Cities" are individuals who moved cities at least once (but never moved countries.) "Moved countries" are those who moved countries at least once. "Not in Dataset" are individuals not used in the estimation (non-movers, or individuals with missing information in either outcome or controls). Standard deviations in parentheses.

Table A2: AKM Variance Decomposition

	Total Sample	**Developing Countries**			**Developed Countries**		
		India	Mexico	Nigeria	US	UK	Italy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: User Log Salaries							
Var(User Effect)	0.234	0.665	0.719	0.713	0.68	0.736	0.784
Var(Firm Effect)	0.671	0.124	0.182	0.16	0.177	0.165	0.231
Var(Covariates)	0.044	0.234	0.122	0.158	0.115	0.13	0.128
Var(Error)	0.049	0.235	0.218	0.229	0.103	0.174	0.163
Cov(User Effect, Firm Effect)	0.035	-0.034	-0.072	-0.064	0.026	-0.036	-0.106
Cov(User Effect, Covariates)	-0.022	-0.107	-0.043	-0.07	-0.059	-0.061	-0.041
Cov(Firm Effect, Covariates)	-0.013	0.012	-0.004	0.004	-0.004	-0.005	-0.005
Var(Log Salary)	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	(0.558)	(0.093)	(0.143)	(0.135)	(0.249)	(0.158)	(0.143)
Panel B: City Mean Log Salaries							
Var(User Effect)	0.019	1.541	0.596	0.313	0.275	0.433	1.204
Var(Firm Effect)	0.899	1.059	0.563	0.559	0.331	0.312	0.75
Var(Covariates)	0.002	0.886	0.244	0.151	0.011	0.051	0.274
Var(Error)	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Cov(User Effect, Firm Effect)	0.06	-0.707	0.069	0.007	0.213	0.166	-0.157
Cov(User Effect, Covariates)	-0.002	-1.027	-0.128	-0.055	-0.011	-0.04	-0.332
Cov(Firm Effect, Covariates)	-0.018	0.491	-0.143	0.036	-0.011	-0.024	-0.125
Var(Log Salary)	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	(0.387)	(0.001)	(0.003)	(0.002)	(0.025)	(0.005)	(0.004)

Notes: Table shows share of explained log salary variance for each part of equation 3. Log salary variance in parentheses. Panel A decomposes the varaince in log salaries across users. Panel B decomposes the variance in average log salaries across cities.

B Bounding Our City Effects Estimates

The bias described in Equation 10 carries through when we aggregate to the city level. Our goal is to construct a conservative lower bound on the share of city-level wage variation explained by true city effects. We begin by defining the true statistic of interest:

$$\Omega = \frac{\text{Cov}(\bar{y}_j, \Gamma_j)}{\text{Var}(\bar{y}_j)}, \tag{14}$$

which captures the share of city-level wage variation explained by true city effects.

In practice, we observe only imputed wages $\hat{y}_{it}$ and compute estimated firm effects $\hat{\gamma}_f$, from which we construct estimated city effects $\hat{\Gamma}_j$. The corresponding analogue is:

$$\hat{\Omega} = \frac{\text{Cov}(\bar{\hat{y}}_j, \hat{\Gamma}_j)}{\text{Var}(\bar{\hat{y}}_j)}. \tag{15}$$

To bridge the gap between these two quantities, we define an intermediate object:

$$\Omega' = \frac{\text{Cov}(\bar{\hat{y}}_j, \Gamma_j)}{\text{Var}(\bar{\hat{y}}_j)}, \tag{16}$$

which captures how much variation in imputed wages is explained by the true city effects.

We assume that true city effects relate to the estimated ones via a linear attenuation relationship:

$$\Gamma_j = \frac{1}{\delta} \cdot \hat{\Gamma}_j + \varepsilon_j, \tag{17}$$

where ε_j is a mean-zero residual uncorrelated with $\hat{\Gamma}_j$, and δ is the ratio of the variance of estimated to true city effects:

$$\delta = \frac{\text{Var}(\hat{\Gamma}_j)}{\text{Var}(\Gamma_j)}. \tag{18}$$

Substituting this relationship into Ω' gives:

$$\text{Cov}(\bar{\hat{y}}_j, \Gamma_j) = \frac{1}{\delta} \cdot \text{Cov}(\bar{\hat{y}}_j, \hat{\Gamma}_j) + \text{Cov}(\bar{\hat{y}}_j, \varepsilon_j). \tag{19}$$

Assuming $\text{Cov}(\bar{\hat{y}}_j, \varepsilon_j) \geq 0$, we obtain a lower bound:

$$\Omega' \geq \frac{1}{\delta} \cdot \hat{\Omega}. \tag{20}$$

This assumption is justified by the presence of positive assortative matching and the structure of the imputation process. Because imputed wages reflect only partial information about worker ability, some of the variation in true city effects is omitted from $\hat{\Gamma}_j$ but remains correlated with $\bar{\hat{y}}_j$. In particular, high-ability workers and high-productivity firms tend to concentrate in cities with higher average imputed wages. This sorting implies that omitted components of the true city effect, captured in ε_j, are likely positively correlated with $\bar{\hat{y}}_j$. Dropping this term thus yields a conservative lower bound on the share of wage variation explained by true city effects.

We now relate Ω' to the true Ω. Let the true average wage in city j be decomposed as:

$$\bar{y}_j = \bar{\hat{y}}_j + \bar{u}_j, \tag{21}$$

where $\bar{u}_j$ is the average imputation error, assumed to be mean-zero and uncorrelated with Γ_j. Then,

$$\mathrm{Cov}(\bar{y}_j, \Gamma_j) = \mathrm{Cov}(\bar{\hat{y}}_j, \Gamma_j), \tag{22}$$

while the variance of true average wages satisfies:

$$\mathrm{Var}(\bar{y}_j) = \mathrm{Var}(\bar{\hat{y}}_j) + \mathrm{Var}(\bar{u}_j). \tag{23}$$

Using the individual-level fit $R^2_{\text{ind}} = \mathrm{Corr}^2(y_{it}, \hat{y}_{it})$, we can conservatively bound:

$$\frac{\mathrm{Var}(\bar{\hat{y}}_j)}{\mathrm{Var}(\bar{y}_j)} \geq R^2_{\text{ind}} \quad \Rightarrow \quad \mathrm{Var}(\bar{y}_j) \leq \frac{\mathrm{Var}(\bar{\hat{y}}_j)}{R^2_{\text{ind}}}. \tag{24}$$

Hence,

$$\Omega = \frac{\mathrm{Cov}(\bar{y}_j, \Gamma_j)}{\mathrm{Var}(\bar{y}_j)} = \frac{\mathrm{Cov}(\bar{\hat{y}}_j, \Gamma_j)}{\mathrm{Var}(\bar{y}_j)} \geq R^2_{\text{ind}} \cdot \Omega'. \tag{25}$$

Combining this with the earlier bound on Ω' yields our final conservative lower bound:

$$\Omega \geq \frac{R^2_{\text{ind}} \cdot \hat{\Omega}}{\delta} \tag{26}$$