The Global Value of Cities

Aakash Bhalothia — UC San Diego
Gavin Engelstad — Northwestern
Gaurav Khanna — UC San Diego*
Harrison Mitchell — UC San Diego

June, 2026

Abstract

What is the economic value of each city around the world, and what can city effects reveal about productivity, migration, and the allocation of workers across space? We study these questions using a new global dataset of job histories for 513 million workers across 220,000 locations in 191 countries. Leveraging worker moves, we implement an event-study design to estimate city earnings premia, separating place-based wage differences from worker sorting. We find that place matters substantially. Within countries, estimated city premia account for 45–73% of wage differences across cities. These estimates reveal large global disparities in the value of place. High-premium cities are larger, more industrially diverse, and allocate more workers to high-premium firms. We use these estimates to quantify the gains from reallocation. The dispersion of city effects is substantially larger in low-income countries, implying significant unrealized gains from migration. Reallocating workers across cities and firms to match the US distribution yields meaningful income gains in developing economies. Together, our results highlight the central role of location in shaping global income differences and point to spatial and within-city allocation as key barriers to development.

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

*Corresponding author: Khanna (UCSD and NBER). Email: gakhanna@ucsd.edu. Website: econgaurav.com
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, Patrick Kline, Miguel Zerecero, Vedant Vohra, Nicolas de Roux, various seminar participants worldwide, and colleagues for their insightful comments. We use Refine.Ink to check the text for consistency and clarity, but no AI was used for data analysis.

1 Introduction

Why are some cities more valuable than others, and what can their value tell us about productivity, migration, and the allocation of workers across space? Answering these questions requires measuring the economic value of every city around the world—estimating, for example, how earnings change when a software engineer moves from Bangalore to San Francisco. Yet observed wage differences across cities may reflect the sorting of more skilled workers into more desirable locations, rather than the value of place itself. Separating place effects from worker sorting 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 leverages 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: how much of global city wage differences reflects worker sorting (Young, 2013; Behrens et al., 2014; Combes et al., 2008), and how much reflects city effects (Ciccone and Hall, 1996; Diamond, 2016; Glaeser and Gottlieb, 2009; Card et al., 2025)? With our estimated city effects, we then study what distinguishes high-value cities and quantify the potential gains from moving workers toward places with larger city effects.

We first distinguish between city effects and individual-level sorting. City effects capture the earnings premium associated with a location, through firms, industries, and agglomeration, while sorting reflects the selection of high-ability workers into high-wage cities. Distinguishing between these mechanisms is essential for using city wage differences to study migration, urban productivity, and aggregate income. If city effects account for a large share of wage differences, moving workers from low- to high-value cities may raise earnings and, under standard wage-productivity assumptions, aggregate output (McKenzie et al., 2010; Clemens, 2011; Clemens et al., 2019). 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 movers across cities and measures changes in earnings, job seniority, occupation, and industry. This approach quantifies how much of the wage differential between city pairs is reflected in within-worker earnings changes around moves. We also account for heterogeneity in firm premia within cities by estimating establishment effects worldwide and aggregating them to the city level.

These estimated city effects give us a common global metric to pursue two main objectives. First, we seek to understand what makes a city more valuable. We estimate establishment effects for firms worldwide and show that high-value cities allocate a larger share of workers to high-premium firms. We relate city effects to observable characteristics—such as economic structure, amenities, industrial composition, and firm diversity—to describe the features associated with high-value cities. Our second objective is to quantify how the allocation of workers across firms and locations shapes potential earnings. The dispersion of city effects within a country measures the potential earnings gains from reallocating workers toward higher-value locations, and we show that these gains evolve over the course of development. In high-income countries, internal migration more frequently directs workers toward high-value cities. Our quantification indicates that reallocating workers both *within* and *across* cities could generate substantial earnings gains in low-income economies under fixed-premia accounting assumptions.

The analysis uses global data on 513 million LinkedIn users and tracks their job histories across 220,000 locations worldwide. The data include job locations, titles, firm and establishment names, and the start and end dates of each work spell. These histories allow us to construct individual-level panels that follow workers across jobs, firms, and cities. Because LinkedIn overrepresents highly educated and white-collar workers, we interpret the analysis as most informative about relatively high-skilled, formal-sector labor markets. Predicted wages are constructed from 200 million salary observations linked to location, job title, tenure, seniority, and company, and are adjusted for purchasing power parity. A key feature of the data is granular job-location information, regardless of where a firm is headquartered. This allows us to track moves across cities worldwide, along with changes in job status, titles, seniority, type of work, and predicted salaries. We use this information to estimate how earnings and job characteristics change when workers move across cities.

Empirically, we begin by documenting stark geographic heterogeneity in city-level wages across the world. These reflect regional incomes (city-level wages are lower in Africa and South Asia than in Europe), country borders (there are sharp jumps at the US-Mexico border), and within-country regions (coastal USA has higher-wage cities than the Midwest).

Our event-study analysis asks whether workers' earnings and job characteristics change around moves between city pairs. Conditional on individual, time, and time-since-move fixed effects, we examine whether moves to cities with higher wages, more senior positions, higher-paying occupations, and higher-wage industries are associated with changes in an individual's earnings, job seniority, occupation, and industry.[1] That is, our event studies show sharp improvements in job quality and earnings when individuals move to higher-value cities.

The event-study design offers several advantages. By conditioning on individual fixed effects, we absorb time-invariant worker characteristics, and the event-study framework allows us to test for pre-trends in wages prior to a move, providing evidence that moves are unlikely to be driven by preceding wage shocks. Identification relies on the assumption that unobserved wage shocks or worker-city match components that occur at the time of relocation are not systematically related to the estimated gains from moving between cities; we provide supporting evidence for this assumption. 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-value cities may attract such workers. As in the broader mover-design literature, our interpretation relies on exogenous mobility and additive separability: moves should not be 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

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

moving between city pairs, supporting these identifying assumptions.

Yet, standard Abowd et al. (1999) AKM frameworks may yield difficult-to-interpret city effects when firm wage premia vary 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 earnings premia. 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 worker quality to firm or city premia. We develop an econometric model that bounds this bias and use matched employer–employee data from Italy to discipline the magnitude of the correction.[2]

The resulting city effects reveal a large role for place, but not an exclusive one. Within countries, city effects account for about 45–73% of wage differentials between cities.[3] Cross-border moves imply an even larger role for place, consistent with randomized visa evidence on large wage gains from moving between India and the US (Clemens, 2013; Clemens et al., 2019), though we interpret these estimates more cautiously because international moves may involve stronger selection, country-level institutions, and match-specific shocks. Overall, country effects are important, but city effects are far from negligible. Moreover, the share explained by city effects varies across countries, with important implications for the gains from internal migration.

Having estimated city effects across the world, we turn to two main objectives: first, identifying the observable features associated with larger city effects. Second, what are the potential aggregate earnings gains from reallocating workers *within* cities, across firms, and *across* cities, within countries? If we additionally make the standard assumption that wages are correlated with productivity and that firm markdowns do not vary systematically across cities, then these earnings gains can also be interpreted as potential output gains.

We answer these questions through several new empirical facts. First, the allocation of workers across firms *within* cities matters. We decompose city effects into (i) the presence of high-paying firms and (ii) the share of workers employed in those firms. Building on Olley and Pakes (1996), we show that worker allocation across the firm-premium distribution is an important component of city effects. In high-income countries such as the US and the UK, a larger share of workers is concentrated in high-premium firms. In contrast, in lower-income countries such as India and Mexico, lower-value cities have fewer workers in high-premium firms.

High-value cities also differ systematically in their observable urban structure. Cities with greater industrial diversity and complexity exhibit larger city effects, while more specialized cities with concentrated industry structures have lower city effects. City effects also increase with city

[2]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).

[3]In the US, 50.8% of wage differences reflect city earnings premia, closely mirroring the 50% finding from Card et al. (2025) on US commuting zones using LEHD data.

size, consistent with agglomeration economies, and with the prevalence of skilled and senior-level jobs (Hornbeck and Keniston, 2017; Greenstone et al., 2010; Yi and Moretti, 2024). Together, these patterns show how estimated city effects line up with urban structure and speak to a long-standing question in urban economics: what distinguishes high-value cities?

City premia measure earnings differences, not welfare. The gap between earnings premia and welfare may vary by development, 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; Rosenthal-Kay, 2026). We examine how city effects relate to amenities, disamenities, and proximity to high-value cities. Pollution is positively correlated with city premia, with a steeper relationship in developed economies, consistent with pollution partly proxying for production intensity or congestion. Climate and consumption amenities, by contrast, are correlated with city premia in ways consistent with compensating differentials, although other differences could generate similar patterns. We also document spatial decay in city premia around high-premium cities, consistent with agglomeration shadows (Hornbeck et al., 2025).

Our next set of results addresses the 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 with moves *within* countries. Potential earnings gains from spatial reallocation rise with migration distance: moves between more distant cities tend to involve larger differences in city effects between origins and destinations. Consistent with this, actual movers tend to relocate over longer distances. 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 highest-premium cities. Many moves in low-income economies are directed toward cities that are not at the top of the city-effect distribution, implying sizable unrealized gains from migration patterns more closely aligned with estimated city effects.

To understand where internal migration could generate the largest gains, we examine how dispersed city effects are within each country. 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 internal migration intensity and weaker business environments, particularly where firm entry is costly and basic infrastructure such as reliable electricity is lacking. We interpret these patterns as suggestive of frictions operating on multiple margins at once: housing supply and land-use constraints in high-value cities that limit migration (Harari and Wong, 2025b), and firm-side barriers that limit firm entry and scale. We do not directly identify which friction binds most tightly, but the correlations suggest that the gap between actual and potential spatial allocation is widest where mobility costs and firm frictions are most likely to bind.

Finally, drawing on our estimated city effects and decomposition results, we conduct a set of partial-equilibrium accounting exercises. We quantify how much average wages could rise if workers were reallocated toward higher-value cities and high-premium firms, holding estimated effects fixed.

First, we reassign each country's population *across* its cities so that the distribution of workers 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 workers in US cities are more concentrated in high-premium firms, we apply the US pattern of *within*-city allocative effects to other countries. This raises average wages by 2.6 percent in India, though the gains are smaller in Mexico and Nigeria, and especially modest in richer countries. Finally, combining both exercises, first shifting workers across firms and then reallocating workers across cities, yields roughly additive effects: 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 exercises do not crowd each other out and that substantial productivity gains remain locked behind both spatial and within-city frictions.

These findings contribute to recent debates on the value of locations for earnings and intergenerational mobility (Card et al., 2026). A recent starting point is Card et al. (2025), who use AKM methods to estimate location premia within the US.[4] The distinctly global nature of our analysis creates certain advantages. First, it allows us to study within-country and cross-country moves in a unified framework. Within this unified framework, cross-border moves are associated with much larger wage gains, though we interpret them more cautiously. This complements work on migration that has focused either on internal migration within countries (Bryan and Morten, 2019; de La Roca and Puga, 2017; Card et al., 2025) or cross-border and country-level place premia (Clemens, 2013; Clemens et al., 2019; McKenzie et al., 2010; Amanzadeh et al., 2024). Second, our global analysis shows that potential gains from reallocating workers across firms and locations vary by country. This connects to work on gains from reallocating inputs across firms and sectors (Hsieh and Klenow, 2009; Gollin et al., 2014). Our methods differ from much of the earlier urban literature: by combining a mover design with event-study diagnostics, we can test for pre-trends that would be consistent with violations of the identifying assumptions, as in Finkelstein et al. (2016) and Badinski et al. (2023).

Next, we speak to fundamental urban economics questions and unpack correlates of these city effects, providing new evidence on what distinguishes high-value cities. Previous work highlights the importance of agglomeration economies (Ciccone and Hall, 1996; Rosenthal and Strange, 2004; Yi and Moretti, 2024), workforce skill composition (Glaeser and Maré, 2001; Moretti, 2004; Diamond, 2016), and urban redevelopment (Gechter and Tsivanidis, 2023), often within particular countries. We provide global evidence on these patterns and how they vary over development (Harari, 2020; Asher et al., 2025; Harari and Wong, 2025b; Harari, 2024; Morten and Oliveira, 2024; Donovan et al., 2023). For instance, population density and skill concentration are more strongly associated with city premia in developing countries. In low-income countries, certain cities have smaller estimated effects partly because fewer workers are employed in high-premium firms. Reallocating workers

[4] 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).

toward higher-premium firms within these cities may raise measured city premia.

Third, we speak to the debate on whether disparities in city wages reflect ability-biased sorting (Young, 2013; Behrens et al., 2014; Combes et al., 2008) or place-based earnings premia (Ciccone and Hall, 1996; Diamond, 2016; Glaeser and Gottlieb, 2009; Card et al., 2025). We find a meaningful role for ability-based sorting: in some countries, city effects explain about 45% of wage differentials across cities. Yet place effects remain economically large and are strongly correlated with economic structure and firm allocation. Our results reconcile these views by showing that sorting and place effects coexist, with their relative importance varying across countries and levels of development.

Finally, we speak to the literature on the gains from migration and the value of place (Clemens et al., 2019; Heise and Porzio, 2021; Clemens, 2011; McKenzie et al., 2010; Khanna et al., 2025). We show that these potential gains are particularly large in low-income countries, where city effects are more dispersed (Morten and Oliveira, 2024). Reallocating workers toward high-value urban centers has the potential to raise aggregate earnings in low-income countries.

Our paper is organized as follows. Section 2 describes the data and the construction of our estimating sample. Section 3 documents descriptive facts about wage differences across cities, migration patterns across city types, and event-study evidence around moves. Section 4 describes our empirical strategy for estimating city effects and discusses potential sources of bias. Section 5 presents the estimated city effects and their implications for heterogeneity in gains from different types of moves. Sections 6–8 document new facts and implications of these city effects. Section 6 documents correlates of city effects, including industrial structure, agglomeration shadows, spatial decay, and the allocation of workers across firms. Section 7 documents the distribution of city effects and potential gains from migration over development. Section 8 quantifies accounting gains from reallocation within and across cities, and Section 9 concludes.

2 Data

Our primary data come from Revelio Labs, which compiles detailed profiles from public professional networking websites, including LinkedIn. Although the data also include harmonized profiles from other sites, LinkedIn accounts for the large majority of users, so we focus on LinkedIn throughout the discussion. The data contain over 513 million professional profiles, created when users enter their education and employment histories. We primarily use the employment histories, which report job location, job title, firm name, and the start and end dates of each position. Using these start and end dates, we construct user-level panels that follow workers across jobs, firms, and locations over their observed work histories.[5]

LinkedIn users differ from the general population. Although profiles cover a wide range of careers, highly educated and white-collar workers are overrepresented. We interpret our estimates

[5] There is a growing literature using online career and wage data, including Glassdoor, to study labor-market outcomes (Martellini et al., 2024).

as informative about the spatial earnings dynamics of relatively high-skilled, formal-sector workers.

Additionally, countries exhibit varying degrees of LinkedIn use, even among college-educated individuals. As noted by Amanzadeh et al. (2024), the ratio of college-educated individuals to LinkedIn users varies across countries, with particularly low LinkedIn use in Central Asia. This variation matters most for cross-country comparisons; most of our analysis estimates city effects within countries. Coverage also varies with income: Appendix Figure A1 plots the country coverage rate from Amanzadeh et al. (2024) against GDP per capita and shows a moderate positive relationship. For this reason, our cross-country analyses control for coverage rates; doing so does not materially affect the estimates.

Seniority, Occupation, and Industry Scores. We measure variation in job characteristics across space using three outcomes: job seniority, occupation scores, and industry scores. The seniority metric is created using an ensemble machine-learning model and converted into an ordinal classification with five categories.[6] In addition to seniority, we construct occupation and industry scores for each employment spell to capture variation in the type of work performed and the 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 leave-one-out average predicted wage among observations with the same O*NET code. Similarly, we define an 'Industry Score' as the leave-one-out average predicted wage across spells in the same NAICS-defined industry. These scores provide wage-relevant benchmarks for the type of work and sector associated with each job.

Wages. We measure earnings using supplementary imputed wage data from Revelio Labs. For each job position on a profile, we observe imputed salaries. Wages are estimated using job titles grouped into 1500 categories, company-specific information, local economic variables such as median housing values and unemployment rates,[7] position-specific information such as tenure and seniority, and company identifiers. The imputation model is trained on over 200 million salaries from a wide variety of sources. To generate global wage predictions, the provider adjusts predictions across countries using purchasing power parity and cross-country differences in wages by job category

[6]The model combines three inputs. First, information about an individual's current job, including title, company, and industry, is used to generate an initial seniority score. Second, details from the individual's employment history, including the duration of prior positions and the seniority of those roles, are incorporated to produce a second score. Third, age is used to create a third score. These three components are averaged into a continuous measure of seniority, which the data provider maps to the most likely category using prediction samples associated with recognizable keywords (e.g., 'junior', 'senior', 'director'). The resulting categories are 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).

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

relative to the US. Wages are also adjusted for country-specific inflation. The imputed wages behave similarly to observed wages from job postings and survey data Amanzadeh et al. (2024).

Locations. A key feature of the data is granular location information for individual positions. When users create a position entry, they report the location of that job. We use this information to match employment spells to cities around the world. This job-level location information lets us identify where individuals work, even when they are employed by multi-establishment or multinational firms. Our global coverage is extensive, with over 220,000 locations in the raw data. Figure 1a plots average salaries for each location in the dataset with over 100 users. Because global income differences can obscure within-country variation, Figures 2a and 2b show city-level coverage separately for the US and India. These maps show dense city-level coverage in both countries and recover familiar regional wage patterns, such as higher wages in California and New England in the US. Amanzadeh et al. (2024) painstakingly validate the coverage of the dataset across countries, and show how well the data matches migration (and return migration) flows.[8]

Analysis Sample. We construct the analysis sample in several steps. For positions with no listed end date, which appear as "present" on LinkedIn, we assign an end date equal to the Revelio Labs collection date at the end of 2023. We then expand the data from the user-position level to the user-year level using each position's start and end dates. We apply a 3% annual wage growth rate backward from the end date of each position, roughly matching average year-on-year salary growth in the US. As we show in an extensive set of robustness checks, our results are robust to also simply using a constant wage or log interpolation between the estimated start and end salaries. Because identification comes from moves across cities, we restrict the main sample to users observed in exactly two cities during the sample period.[9]

Additional Data. We supplement the LinkedIn data with several external data. We obtain city-level population and density measures from GeoNames and the Global Human Settlement Layer Functional Urban Areas (GHS-FUA). We construct measures of pollution ($PM_{2.5}$), temperature, and precipitation over a 20-year period using data from Shen et al. (2024) and Harris et al. (2020). For the US, we use Zillow home price data to test whether adjusting for housing costs changes our estimates. At the country level, we use the World Bank's World Development Indicators and B-READY data to study how the role of city effects varies with country characteristics. We also digitize data from Bell et al. (2015) to obtain a harmonized measure of internal migration intensity. Finally, we assess the importance of wage imputation by repeating our city-effect estimation using

[8] In order to increase our confidence in the reliability of the data, we drop countries with extremely low coverage rates, very few unique locations, or with potential problems with how locations were recorded in the data. The most notable countries we drop through this process are China, Japan, and Brazil.

[9] The main sample excludes individuals who worked in more than two cities during their careers to avoid concerns about return migration or attributing wage gains to the wrong location based on move order. Table A4 shows that the results are similar when we use all moves.

matched employer-employee data from Italy, where true wages are observed.

3 Descriptive Facts

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

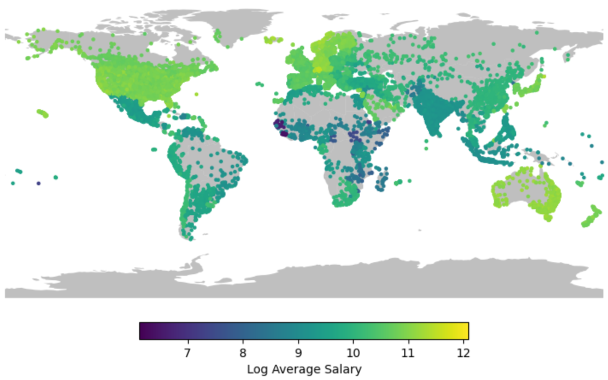

(a) Average Log City Wages, Global

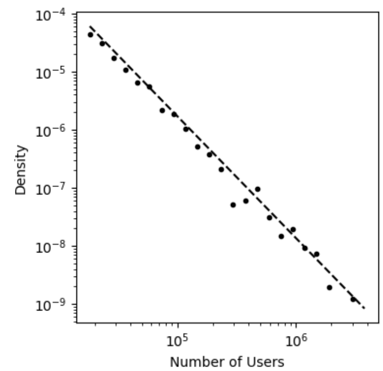

(b) Cities by Number of Users

Notes: The left panel shows the average wages calculated in each city with over 100 users in our data (on a log scale), across the world. The right 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.

We begin with descriptive facts about the city distribution in our data. Figure 1a plots average city-level wages around the world. Three patterns stand out. First, the data provide broad geographic coverage. The blank grey zones largely correspond to sparsely populated areas. Second, average wages vary sharply across regions, with lower wages concentrated in Africa and South Asia and higher wages concentrated in Europe and the US. Third, country borders are starkly visible in some regions, suggesting large cross-border wage gaps. This is particularly evident along the Mexico-US border and between North Africa and Southern Europe.

The distribution of users across cities follows a familiar urban hierarchy. Figure 1b shows that users are concentrated in a small number of major cities, while most cities have far fewer users, consistent with a power-law. This pattern mirrors the global distribution of city populations.

To better show within-country variation, Figure 2 maps average city wages separately for the US and India, allowing the color scale to vary by country. The maps highlight both regional wage variation within countries and the density of city coverage, even after restricting the sample to cities with more than 100 users. In the US, wages are higher on the coasts than in much of the Midwest, though several Midwestern cities also have high salaries. In India, high-wage cities are more common in the south and west, while lower-wage cities are more common in the center and east. Appendix Figure A2 shows analogous maps for the UK and Italy: in the UK, high-wage cities are concentrated in the south, where coverage is also denser; in Italy, the highest-wage cities cluster

Figure 2: Average City Wages in the US and India

(a) USA

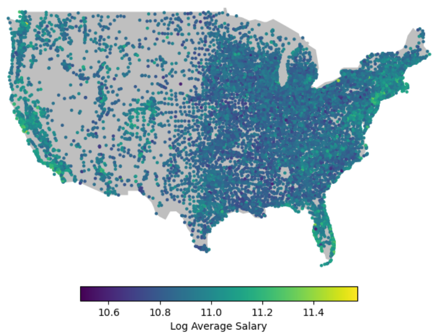

(b) India

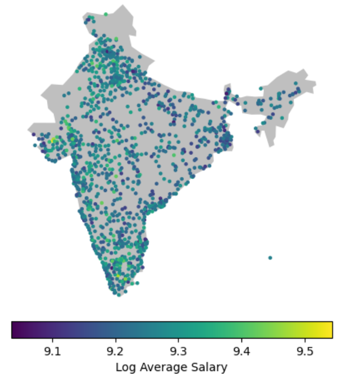

Notes: The maps plot average log(salaries) in each city. Legend scale varies by country. Sample restricted to cities with more than 100 LinkedIn users.

around Lombardy and the north, while wages are lower in the south. Importantly, the wage range on the log scale varies substantially across countries.

3.1 Transition Matrices and The Movement of Workers Across City Wages

Our estimation relies on workers who move across firms. We therefore describe both the mover sample and the moves they make. Appendix Table A1 compares individuals who move across firms within cities, move across cities within countries, move across countries, and remain outside the estimation sample. Mean salaries may differ across columns for two reasons: different types of people move, and moving itself may change outcomes. Our regressions address the first concern using individual fixed effects; the second is the primary object of our study. Despite these possibilities, the differences across columns are small and not statistically significant. Within countries such as the US and India, mean salaries are also similar across groups. These patterns suggest that movers and non-movers are similar on average, but our estimation does not rely on that assumption. Our empirical design allows movers to differ systematically from non-movers.

We next use the mover sample to describe the moves themselves. Figure 3a shows the distribution of movers by origin income. We divide cities into five average-wage quintiles to summarize transitions across the city wage distribution. Moves are most common from the middle of the city wage distribution. This pattern is consistent with migration models in which out-migration is highest from middle-income places: low-income individuals may face binding migration costs, while high-income individuals have weaker incentives to leave.

Figure 3b then shows the transition matrix for all recorded cross-city moves, with origin-city quintiles on the horizontal axis and destination-city quintiles on the vertical axis. Most moves occur along the diagonal. This likely reflects that most moves occur within countries and that country

Figure 3: Movers Sample: Income Quintiles and Transition Matrices

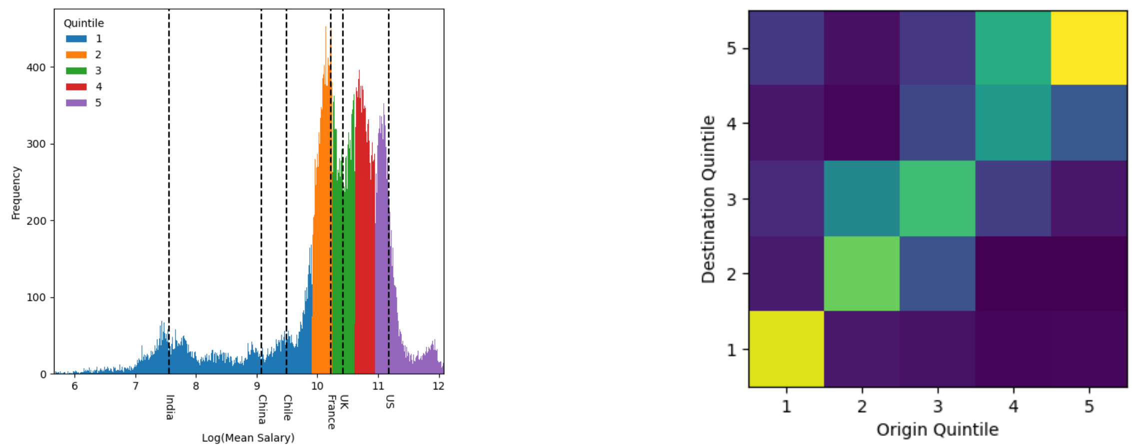

(a) Wage Distribution of Locations by Movers

(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.

borders account for large differences in average city wages. There is also visible mobility from the fourth origin quintile to the fifth destination quintile, suggesting relatively high upward mobility from upper-middle-wage origin cities.

Finally, in Figure 4, we examine transitions within and across borders. Panel A focuses on within-country moves. Higher-wage quintiles attract a larger share of internal migrants. This, reassuringly, suggests potential gains to aggregate incomes driven by migration. In low-income countries, some migrants move to fourth-quintile cities, while in richer countries most migrants move to top-quintile cities. The brighter northwest triangle shows that many moves are from lower-wage cities to higher-wage cities. The dense top-right cell shows that workers from top-quintile origin cities often move to other top-quintile cities and rarely move to bottom-quintile cities. In India, most movers go to cities in the top two quintiles, regardless of their origin quintile.

Panel B turns to cross-border moves. The axes report income quintiles within origin and destination countries.[10] Here, patterns are even starker. In low-income countries, cross-border migrants are concentrated in the top two destination quintiles; in high-income countries, they go to the highest-wage cities. This reflects the notion that cross-border migrants may improve allocation (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). The US stands out: international migrants are especially concentrated in high-wage destinations.

Together, these patterns show that movers tend to move toward higher-wage cities. The next step is to distinguish whether high-wage cities pay more because higher-earning workers sort into

[10] 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.

Figure 4: Transition Matrices for Migrants

Panel A: Internal Moves

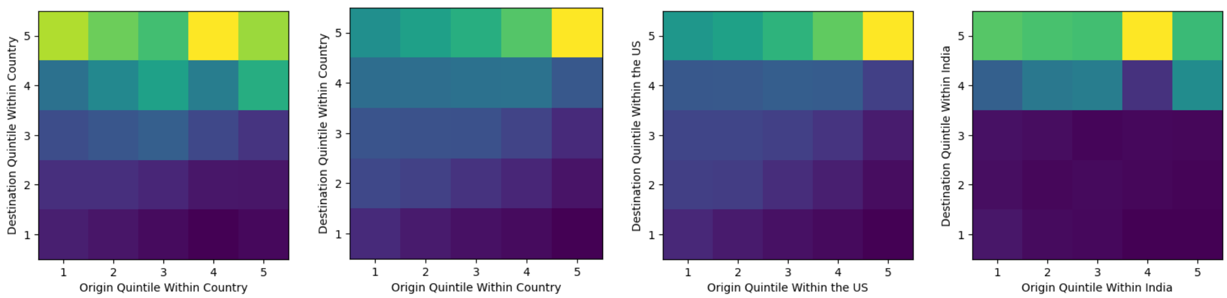

(a) Low Income Countries (b) High Income Countries (c) In the US (d) In India

Panel B: Cross Border Moves

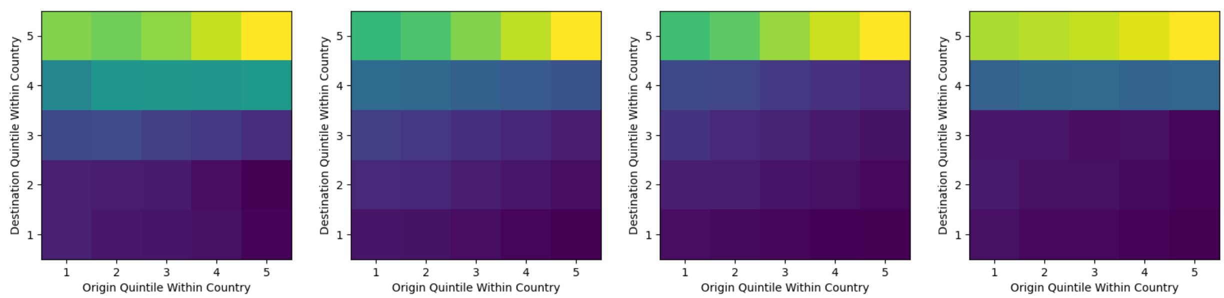

(e) Into Low Income (f) Into High Income (g) Into the US (h) Into India

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).

them or because workers' earnings rise after they move there.

3.2 Descriptive Event Studies

We begin by studying bilateral moves between pairs of cities in an event-study framework. If wage differences across cities are primarily driven by sorting, movers need not experience systematic wage changes when they change cities. Conversely, if cities have place-based earnings premia, movers to cities with higher average wages for other workers should experience wage gains. In comparison, those moving to cities with lower average wages would face wage losses.

To investigate this empirically, we define δ_i as the difference in average job characteristics (wages, seniority, occupation score, or industry score) between mover i's destination and origin cities. Specifically, for mover i with origin city $o(i)$ and destination city $d(i)$:

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

where $\bar{y}_j$ is the average of $\bar{y}_{jt}$ across years, and $\bar{y}_{jt}$ is the mean of outcome y_{it} among workers in city j in year t. Following Finkelstein et al. (2016), we estimate:

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

where y_{it} is the outcome for individual i in calendar year t: log wage, job seniority, log occupation score, or log industry score. The individual fixed effect α_i captures time-invariant worker characteristics, and τ_t controls for calendar-year fixed effects. $I_{r(i,t)}$ is a vector of relative-year indicators, where $r(i,t) = t - t^*$ for mover i who moves in year t^*. The relative-year coefficients $\theta_{\mathrm{r}(i,t)}$ are the main parameters of interest. They capture changes in y_{it} around the move, scaled by the destination-origin difference δ_i. When the outcome is earnings, δ_i is the difference in average wages between the destination and origin cities, and y_{it} is individual i's wage.

We present estimates from Equation 2 for different mover populations. On the y-axis, we plot estimated coefficients $\theta_{r(i,t)}$ for the corresponding relative year $r(i,t)$ on the x-axis. The coefficients capture changes in job outcomes around the move, scaled by the average wage difference between the mover's destination and origin cities. We normalize the value for $r(i,t)$ = -1 to 0.

Panel A of Figure 5 plots estimated coefficients $\theta_{r(i,t)}$ for the full mover sample. When the outcome is log wages, the figure shows a sharp jump in the year of the move, from zero to approximately 0.95; this estimate is driven largely by cross-border moves. This implies a city-effect share of $\approx 95\%$ in observed wage differences across cities globally (Finkelstein et al., 2016), though this global estimate is driven largely by cross-border moves. For other outcomes, such as seniority, occupation score, and industry score, the estimated jumps explain a smaller share of cross-city differences. This suggests that perhaps technological differences across countries may account for large wage gaps without generating equally large differences in observed job characteristics.

Figure 5 Panel B, presents results for those who move to cities within the same country. For log wages, the jump of 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 123,431 cities, while 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 presents log-salary event-study estimates for selected developed and developing countries. Some post-move trends suggest that δ_i is positively correlated with wage growth after the move. Appendix Figures A4 and A5 report the corresponding estimates for seniority, occupation score, and industry score. For a few emerging economies, the industry-score effects are zero or negative, suggesting that movers to higher-wage cities may end up in lower-scoring industries. This may

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

Panel A: All Movers

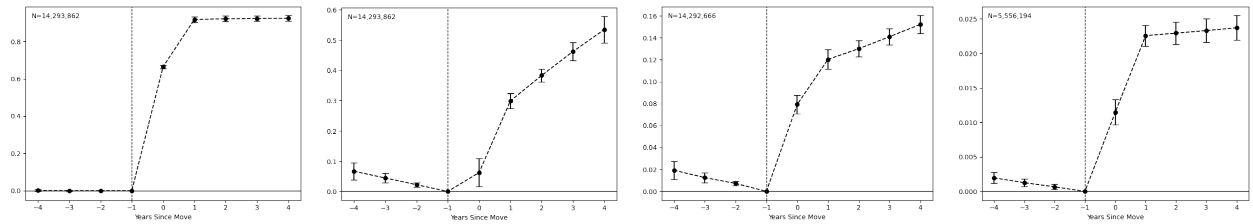

(a) Log Salary (b) Seniority (c) ONET Wage (d) NAICS Wage

Panel B: Within-Country Movers

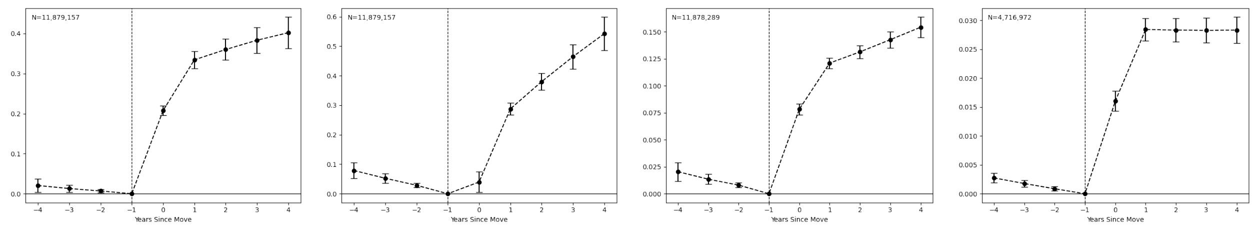

(e) Log Salary (f) Seniority (g) ONET Wage (h) NAICS Wage

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.

reflect frictions in transferring skills across industries or other job-search frictions in low-income settings. Our decomposition in Section 4.1 allows for this possibility by separating city effects from the allocation of workers across firms.

Appendix Figures A3–A5 assess the robustness of the event-study results across numerous countries and types of outcomes. These figures show that the patterns are mostly unchanged when we interpret salary growth in different ways. For instance, one concern is that assumptions about salary growth during non-reported years within a position could affect the trends. The appendix figures show that they do not. Finally, Appendix Figure A7 shows that our estimated city effects are robust to controlling for differences in local prices, specifically house prices. Using Zillow home price data and only our US sample, we deflate wages by dividing them by the average home value within each city-year; when we re-estimate city effects in the US, the adjusted and unadjusted estimates have a nearly one-to-one relationship.[11]

3.3 Average Gains by Origin-Destination Pair

The event study graphs show stark increases in earnings, as a function of the origin-destination wage differential.[12] Keeping with the analysis of bilateral (origin-destination pair) moves, here we

[11]We obtain similar results with a smaller sample when we instead adjust using Zillow rent prices.

[12]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 6: Event Studies for Internal Movers: Log Salary, Selected Countries

Panel A: Developed Economies

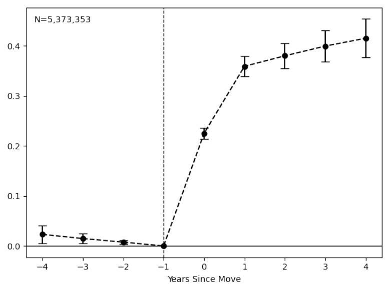

(a) United States

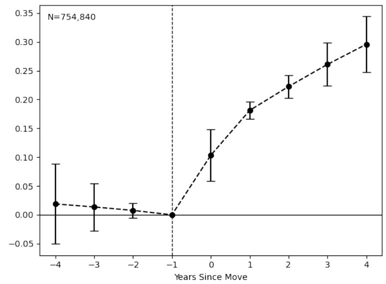

(b) United Kingdom

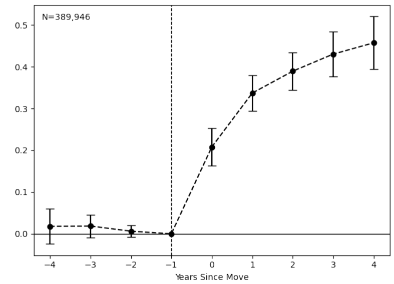

(c) Germany

Panel B: Developing Economies

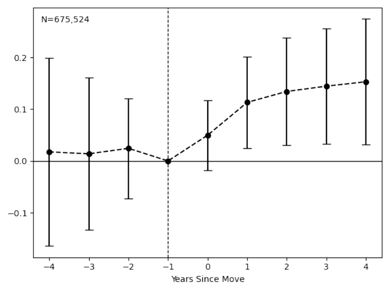

(d) India

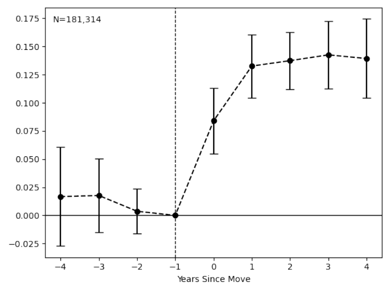

(e) Mexico

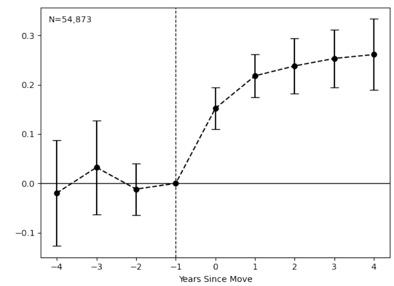

(f) Chile

Each panel presents event study coefficients for log salary among 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.

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 ($\delta_i \equiv \bar{y}_{d(i)} - \bar{y}_{o(i)}$).

This exercise clarifies how the event-study estimates map into city effects. The event-study figures estimate the average wage jump for movers across cities. But this average may mask heterogeneity and nonlinearities across origin-destination wage gaps. For instance, movers between two high-wage cities may experience a different proportional wage gain than movers from a low-wage city to a high-wage city. Figure 7a plots the relationship between the wage jump for each origin-destination pair and the corresponding difference in average city wages. The slope of the fitted line summarizes the average event-study jump across origin-destination wage gaps. We also plot the density of observations along each axis.

The relationship in Figure 7a is roughly linear, though flatter near the center. The flatter center reflects the concentration of within-country moves with smaller wage differentials, while larger differentials mostly correspond to cross-border moves. Because country borders are associated with much larger wage differences, cross-border moves generate larger measured wage changes.

International moves in Figure 7b show a stark, almost one-to-one relationship with average city wage differences. The Bangalore-San Francisco wage differential is mirrored in movers' wage

Figure 7: Change in Wages by Difference in Pairwise City Wages

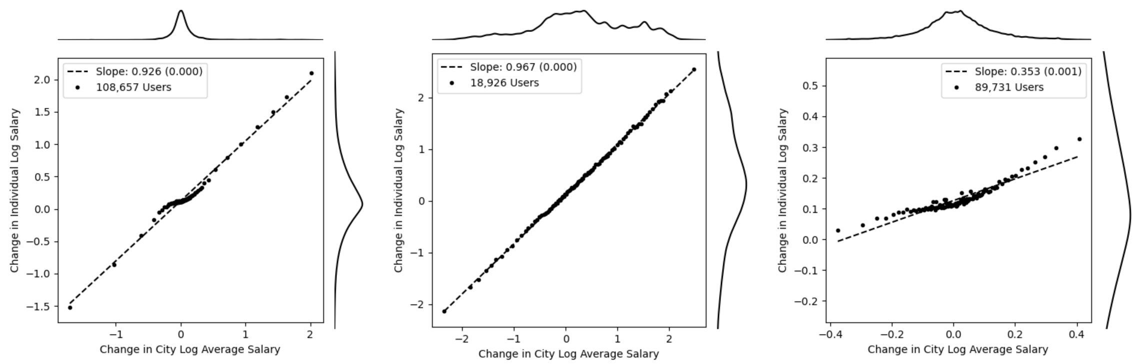

(a) Full Sample (b) International Moves (c) Internal Moves

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.

changes: on average, movers between these cities earn almost the full average wage difference. In contrast, for internal moves (Figure 7c), the relationship is a lot flatter (a slope of 0.34). Movers between Omaha and San Francisco experience a smaller wage increase as a share of the average wage difference between the two cities.

These patterns are consistent with lower migration barriers within the US: high-earning workers may have already sorted across cities, so raw wage gaps partly reflect composition. That is, a fair amount of the wage differential between Omaha and San Francisco reflects ability-based sorting.

We then repeat the exercise within specific countries. Figure 8a shows a slope of 0.37 for the US, similar to the average event-study jump. The slope is much flatter in India: Figure 8d shows a slope of 0.13. The differences in this relationship between Figures 8a-8c and Figures 8d-8f reflect the fact that the changes in salary associated with moving are typically higher in developed countries compared to developing countries.

We also examine whether the relationship between wage gains and city wage differentials changes over time. Figure A8 plots this relationship by move year. The relationship is stable for international moves, but increases over time for internal moves: between 2010 and 2020, the average wage increase associated with a given origin-destination wage differential rises by nearly 35%. This pattern could reflect a growing role for city effects, rising frictions such as congestion or housing supply constraints, or changes in the extent of ability sorting across space.

4 Empirical Strategy to Estimate City Effects

We now move from bilateral origin-destination wage differences to estimating an overall effect for each city. To fix ideas, consider the standard specification from the mover-design literature (Abowd

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

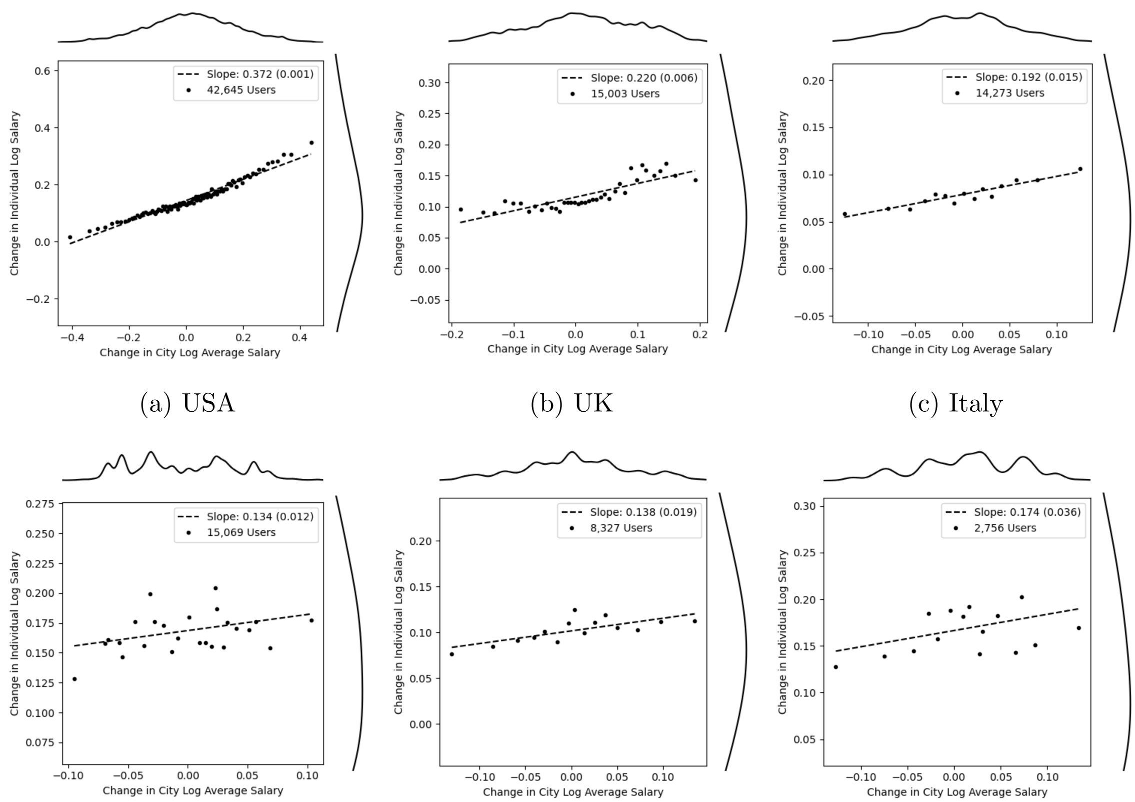

(a) USA (b) UK (c) Italy

(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.

et al., 1999; Finkelstein et al., 2016). We begin with a simple two-way fixed-effects model in which the 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 term 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 is employed in year t. α_i captures time-invariant individual characteristics, and τ_t controls for calendar-year fixed effects. $x'_{it}\beta$ controls for time-varying characteristics, including year relative to the move. These relative-year effects allow the decision to move to be correlated with wage shocks, such as layoffs that lead workers to relocate.

This empirical setup allows for several important forms of non-random mobility. Identification comes from workers moving across cities, so city effects are identified only within the largest connected set of cities linked by worker moves. The model allows mobility to be related to time-invariant worker characteristics and city characteristics. For example, workers may be more likely to move from low-wage or low-amenity cities to high-wage or high-amenity cities. More productive workers may also be more likely to switch cities; individual fixed effects account for this time-invariant selection. Together, these features allow for assortative matching, in which high-earning workers are more likely to move to high-wage cities.

Two standard AKM assumptions remain necessary for interpreting ψ_J as a city effect. First, individual α_i and city effects ψ_J must be log-additively separable, ruling out worker-city interactions. This assumption is consistent with a simple log-linear production function, or multiplicative worker and city effects in levels (Eeckhout and Kircher, 2011; Borovičková and Shimer, 2024).

Second, exogenous mobility requires that, conditional on individual and city fixed effects and other time-varying controls, the error term ϵ is uncorrelated with changes in $\boldsymbol{J}(i,t)$. Following Card et al. (2013), we focus on three potential violations. The first is sorting based on worker-city match quality (Roy, 1951): if such sorting occurs, different workers may receive different wage premia in the same city depending on match quality. While assortative matching does not affect the interpretation of our estimates, sorting on match quality would.

A second violation is drift in individual earnings potential, such as employment shocks or human-capital accumulation, correlated with mobility across cities. Workers who lose jobs or are demoted may be more likely to move. Outside offers from other cities may temporarily raise wages in the current city before a move. Adaptation to a city may raise wages over time, and workers who adapt more quickly may differ in their subsequent mobility. Conversely, workers whose wages fall after moving may be more likely to leave. If moves are driven by one-time signing bonuses rather than permanent wage gains, the event study would show a post-move spike followed by a decline. More generally, mobility correlated with transitory wage fluctuations would violate the design.

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.

The event-study evidence addresses drift and transitory shocks. We find no meaningful pre-trends that would suggest drift in the portable component of workers' earnings power. We also find no 'Ashenfelter's dip' before moves and no one-time spike immediately after moves. Importantly, the event-study framework allows movers to differ from non-movers in time-invariant earnings levels and in wage trajectories around moves.

If mobility reflects idiosyncratic match quality, or if additive separability fails, then moves between two types of cities, say ψ_A and ψ_B, should have asymmetric wage effects. A worker moving from City A to City B would not experience a wage gain that mirrors the wage loss for workers moving from B to A.[13]

We test this symmetry in Figures 7 and 8. We plot average wage changes around moves on

[13]Specifically, for moves from City A to City B:

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

while for moves in the opposite direction:

the y-axis against differences in average city wages on the x-axis. Points to the left of zero capture moves from high-wage to low-wage cities, while points to the right capture moves from low-wage to high-wage cities. Wage changes for moves from low- to high-wage locations are approximately symmetric to those for moves from high- to low-wage locations (Finkelstein et al., 2016). This provides support for additive separability and against large idiosyncratic match components.

Taken together, the absence of pre-trends and the symmetry in wage changes support the use of this model to estimate city effects on log wages.

4.1 Hierarchy Effects

The basic AKM framework may be difficult to interpret when firm wage premia vary within cities. For instance, if movers from low-wage cities leave above-average-paying firms but join below-average-paying firms in high-wage cities, estimated city effects would include what Card et al. (2025) refers to as a 'hierarchy effect.'

We implement the solution proposed by Card et al. (2025). Rather than estimate city effects directly, we estimate establishment effects and aggregate them to the city level. The estimating equation is:

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

where $\boldsymbol{f}(i,t)$ indicates the establishment where individual i is employed in year t. We define the city-specific wage premium as the employment-weighted average of establishment effects within the city:

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

where $j(f)$ maps establishment f to its city, γ_f is the establishment effect, and N_f is the number of person-year observations for establishment f in the estimation sample. Thus, if a worker were randomly assigned to a firm in destination city j rather than origin city j', predicted earnings would differ by $\Gamma_j - \Gamma_{j'}$.

Estimating establishment effects exploits a richer mobility network than relying only on direct city-to-city moves. If movers from Cities B and C enter different establishments in City A, moves between those establishments identify their relative premia. This separates origin-city differences from destination-firm differences, connects cities without direct mover flows, and improves the aggregation of establishment effects into city effects.

This aggregation matters when hierarchy effects are correlated with city effects. A conventional city-level AKM can still be informative even if movers leave high-wage firms and join low-wage

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

where the final expectation term is the bias term. Without this bias term, the effects of moving should be symmetric: $\psi_B - \psi_A$ for moves from A to B and $\psi_A - \psi_B$ for moves from B to A.

firms. But if these firm-premium differences are systematically correlated with city effects, the city-level AKM estimate would not recover the quasi-random assignment thought experiment we describe. Equation 4 also relies on the standard firm-level AKM assumptions, which have been tested extensively in applications across countries (e.g., Card et al. (2013)).[14]

We use this strategy as our main empirical strategy to estimate the city effects in Section 5.

4.2 Decompositions

We decompose wage differences between high- and low-wage areas, focusing on the share of the wage gap between two regions, R and R', 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:

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

The corresponding 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 using 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 use wage estimates imputed by Revelio. Imputed wages are widely used in studies of mobility and historical labor markets.[15] A key limitation is that imputation may compress wage variation because it cannot fully capture individual-specific determinants of earnings.[16] To our knowledge, the implications of wage imputation for AKM estimates have not been formally characterized. We develop a simple model of the imputation process and use it to sign and bound the resulting bias in estimated city effects. We summarize the model and its main implication here;

[14] In practice, we restrict the establishment-level estimation to establishments with a FactSet Entity ID and at least five workers in our data.

[15] Examples include Abramitzky et al. (2016), Abramitzky et al. (2021), and Eli et al. (2016), who use "wage scores" based on average sectoral wages from the 1950 US Census applied backward to full-count census data, and Amanzadeh et al. (2024), who use the same LinkedIn data to study return migration.

[16] Imputation is common in large-scale surveys and censuses. Rates are high even in benchmark US datasets such as the ACS and CPS, where earnings imputation reaches roughly 30% (Bollinger and Hirsch, 2006, 2013).

the full derivation, the bound, and a validation exercise in a setting where we observe true wages using detailed employer-employee matched data are in Appendix B.

Suppose true wages follow:

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

where α_i is a worker fixed effect, $\gamma_{f(i,t)}$ is a firm fixed effect, and ε_{it} is an idiosyncratic shock orthogonal to worker and firm effects.[17] Standard AKM identifies γ_f from within-worker wage changes around moves; α_i differences out regardless of how strongly workers sort across firms. AKM on true wages is therefore unbiased under arbitrary positive assortative matching. The bias we characterize is a property of *imputation*, not of AKM, and arises because the imputation does not observe worker ability directly. We decompose ability into an observable component used by the imputation and a residual,

$$\alpha_i = X_i\beta + \eta_i, \qquad \mathbb{E}[\eta_i \mid X_i] = 0, \tag{10}$$

where X_i collects covariates available to the imputation, such as job title, tenure, and seniority, and η_i is residual ability. The imputation predicts wages from a noisy proxy for $X_i\beta$ together with firm indicators. When workers sort positively across firms, the firm indicators absorb the worker ability that the proxy fails to capture, so the firm component of imputed wages is inflated.

We show in Appendix B that imputed wages load on the firm effect by a factor $(1 + B)$, where

$$B \equiv \underbrace{(1-\lambda)\phi}_{\text{imputation-quality channel}} + \underbrace{\theta}_{\text{PAM-on-unobservables channel}} \geq 0. \tag{11}$$

Here $\lambda \in [0, 1]$ is the reliability of the proxy (the share of true-ability variation it captures), and $\phi, \theta \geq 0$ measure sorting on observable and unobservable ability, respectively (all defined formally in Appendix B). The first channel, $(1 - \lambda)\phi$, vanishes if the imputation uses observables perfectly ($\lambda = 1$) or if there is no sorting on observables ($\phi = 0$). The second, θ, reflects sorting on ability the imputation cannot observe and is not removed by a richer imputation. Both require sorting, so $B = 0$ whenever $\phi = \theta = 0$.

Because time-invariant worker ability and mean-zero imputation noise difference out around moves, AKM on imputed wages identifies firm effects scaled by $(1 + B)$,

$$\tilde{\gamma}_f = (1 + B)\,\gamma_f + \text{const}, \tag{12}$$

which aggregates Equation (5) to city effects scaled by the same factor,

$$\tilde{\Gamma}_j = (1 + B)\,\Gamma_j + \text{const}. \tag{13}$$

[17] We suppress time fixed effects and controls for clarity; they simply carry through the derivation.

If city-mean worker ability is collinear with the city effect (stated formally in Appendix B), the share Ω of city-mean wage variance attributable to city effects (Equation (8)) inherits the same inflation,

$$\hat{\Omega} = (1 + B)\,\Omega. \tag{14}$$

Since $B \geq 0$, imputation-based estimates are upward biased, $\hat{\Omega} \geq \Omega$, and any upper bound on B delivers a conservative lower bound on Ω. Appendix B constructs such a bound using a reverse-regression argument that requires only the standard AKM noise structure.

Validation in Italian matched data. We discipline the magnitude of this bias using matched employer–employee data from Italy, where we observe true wages and can construct imputed wages following the structure of the Revelio procedure.[18] Using true wages, we estimate $\Omega = 0.584$; a detailed imputation designed to follow the global procedure gives $\hat{\Omega} = 0.658$, an upward bias of $B = 0.127$ (7.4 percentage points of the share). Consistent with the model, estimated city effects are more dispersed under imputation than under true wages (Figure A16b). Applying the Appendix B bound recovers a conservative $\Omega = 0.568$, below the true 0.584. We carry this correction we obtain in Appendix B to the global estimates throughout. Appendix B reports the full validation results, the construction details, and a test to assess the role measurement error in the imputation process could play in impacting our results. In order to assess the potential for variation in imputation quality across countries, we add noise to our imputed Italian wage data and find that differential measurement error in the imputation process is unlikely to drive cross-country variation in the importance of the city effects we estimate (Figure A16c).

4.4 Does our Estimation Truly Capture City Effects?

The empirical strategy in the first half of the paper is designed to recover credible city effects; the main contribution is the global scope of these estimates, which allows us to study how the economic role of cities varies over development. We build on established mover-design methods (Card et al., 2025; Finkelstein et al., 2016), but several features of our setting require additional discussion. Having discussed exogenous mobility, we now turn to other concerns about the interpretation of the estimated city effects. We do not have definitive answers to all concerns, but we try to be transparent on possible violations of these assumptions.

First, identification comes from movers, who may differ in important unobserved ways from stayers. For instance, workers with lower moving costs or higher unobserved ability may be more likely to move. Table A1 suggests that movers and non-movers are similar in many respects. And individual fixed effects absorb average ability differences.

Even after accounting for average differences, there may be systematic differences in 'returns'

[18] The dataset was developed by Giuseppe Tattara and the Economics Department at Università Ca' Foscari Venezia. The data cover work histories of Italian workers from 1975 to 2001 and include firm identifiers, locations, firm tenure, and a coarse salary-type measure that proxies for occupation.

to changing cities. Movers may benefit more from relocation, especially if they select destinations where their gains are larger. That is, people with higher expected gains may be more likely to move. This need not bias the mover design, but it affects interpretation: 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 7 and 8 suggest that the effects are similar when moving from Omaha to New York, as is vice versa. A separate advantage of our establishment-level approach is that it uses workers moving across firms *within* cities, in addition to cross-city movers. These within-city moves identify relative establishment effects, expand the connected set of establishments, and allow us to aggregate establishment effects into city effects even when movers from different origin cities enter different firms in the destination city. Finally, Figures 12 and A11 show that residuals are flat with respect to average city wages, providing evidence against unobserved factors that violate exogenous mobility.

Another concern is that, even after accounting for permanent worker differences with individual fixed effects, unobserved time-varying factors may influence both the decision to move and the post-move wage trajectory. Related dynamic concerns are explored in Bonhomme et al. (2019); de La Roca and Puga (2017). These concerns are most directly relevant for variance decompositions rather than mean city effects; nevertheless, all estimates include time-since-move fixed effects, and the event studies examine pre-trends and post-move spikes. For instance, if workers experience positive wage shocks, acquire new skills, or use outside opportunities to move to higher-paying firms, we would expect to see pre-trends in the event studies.

A third set of concerns involves selection into the sample. LinkedIn coverage varies across countries in likely systematic ways. Workers in some locations may be likely to report successes on LinkedIn. Different types of people select into LinkedIn in San Francisco and Lagos. Our estimates are less sensitive to differences in LinkedIn composition because they are identified from within-worker wage changes, with individual fixed effects absorbing permanent differences across users. Remaining concerns about differential coverage are most relevant when comparing estimates across countries and over development; we therefore control for coverage rates in these analyses and find that they do not materially affect the results. Coverage also increases substantially over time, especially after 2006, but our event-study coefficients remain stable. A related concern is that data quality may be lower in developing countries. The noise exercise in Figure A16c addresses this by showing that even substantial degradation in imputation quality leaves our bounds relatively tight.

A fourth set of concerns involves what wages capture, including employer market power and local prices. If firms have market power, movers to a city may have less bargaining power and earn less than their marginal product. Our use of predicted wages may actually help address this, as they are based on seniority, job title, and firm characteristics and therefore less likely to reflect individual-specific bargaining. Within-city moves provide an additional safeguard when market power and bargaining conditions vary across firms. These moves identify relative wage premia across establishments in the same city. Thus, if cross-city movers disproportionately enter firms where bargaining conditions are weaker, those firm-level differences are less likely to be mistaken for city-wide effects. Consistent with this interpretation, Figure 12 shows that earnings changes closely track city effects after accounting for firms' positions on the local job ladder. We further address local prices by re-estimating US city effects after deflating wages using Zillow housing costs. The adjusted and unadjusted estimates have an almost one-to-one relationship (Figure A7).

A related concern is selection into particular origin-destination pairs. For example, workers moving from Houston to Nairobi may be disproportionately likely to be oil executives or government officials and may earn higher salaries. The establishment-effect analysis is designed to address this concern. Even in the city-level event studies, we do not see patterns consistent with this concern: wage changes are positively related to average wage differences, and moves from Houston to Nairobi have effects that are similar in magnitude and opposite in sign to moves from Nairobi to Houston.

A final concern is limited mobility bias, which can arise when estimates rely on cities with very few moves (Bonhomme et al., 2019; Kline et al., 2020; Kline, 2024). This concern is especially important for second moments, as in traditional AKM decompositions. We follow the recommendations in Bonhomme et al. (2019) by relying on cities and firms with substantial mobility and by focusing on first moments, namely 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 9: City Effects, Global

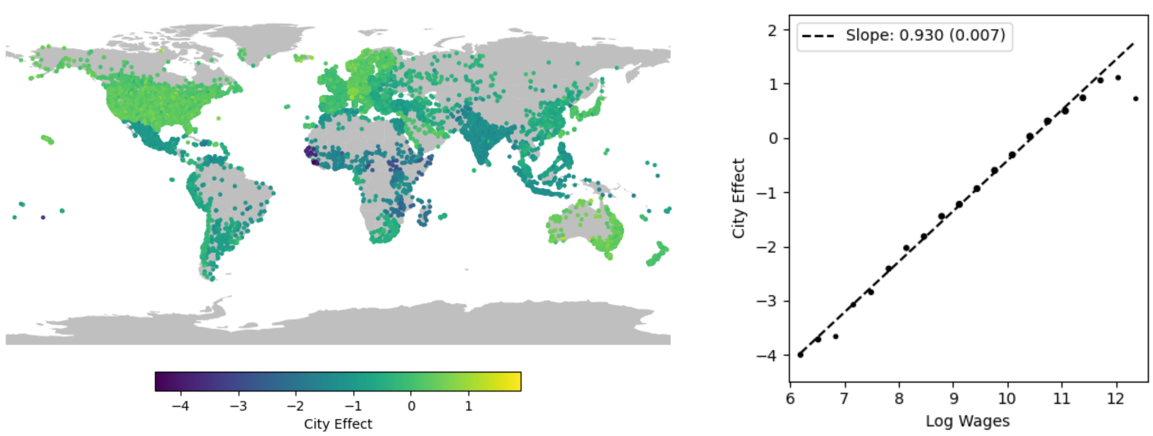

(a) Estimated City Effect (b) Log City Wages vs City Effect

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. $N = 8,025,420$

Using the mover sample and Equations 4 and 5, we estimate city effects for cities worldwide. We use these estimates for three complementary analyses. First, we plot global city effects identified

Figure 10: City Effects, USA

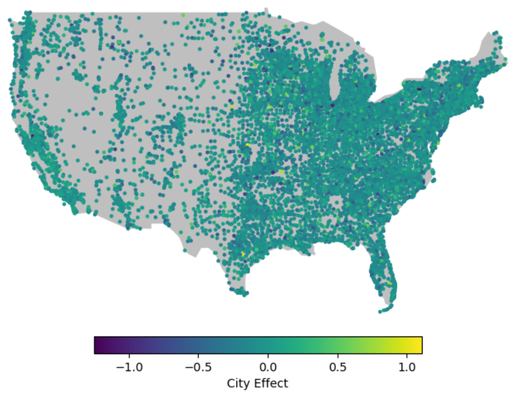

(a) Estimated City Effect

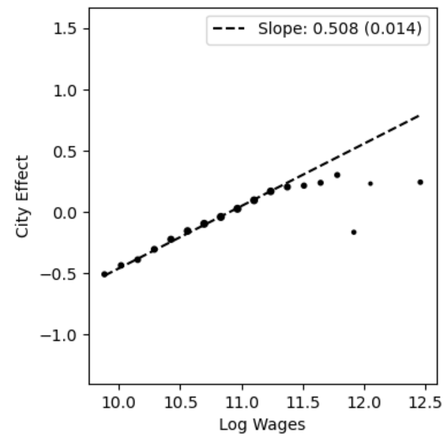

(b) Log City Wages vs City Effect

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.

from international movers (Figure 9) and compare them with estimates based on within-country movers (Figures 10, 11, and A10). This comparison shows how national borders shape measured city effects. Together, these analyses show how the role of city effects changes with migration costs: across countries, across cities within a country, and across national contexts. In subsequent sections, we examine the economic forces associated with these city effects and quantify the potential gains from migration within and across countries.

5.1 City Effects Across and Within Countries

Figure 9a maps the estimated city effects. The map resembles the raw average-wage map in Figure 1a, again highlighting the importance of regions and country borders. Figure 9b plots the relationship between city effects and average city wages. These are strongly positively correlated for cities across the world. The slope of 0.93 indicates that city effects account for most observed wage differences across cities worldwide.

Figures 10 and 11 focus on city effects estimated from within-country moves in the United States and India. The resulting maps differ notably from the raw salary maps 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 10 are more evenly distributed across space. The relationship between city effects and average city wages is flatter within the US ($\beta = 0.51$) and India ($\beta = 0.73$) than in the international comparison ($\beta = 0.93$), indicating that city effects account for a smaller share of within-country wage dispersion than cross-country wage differences.

These findings align with the event-study evidence in Figure 5, where wage changes associated

Figure 11: City Effects, India

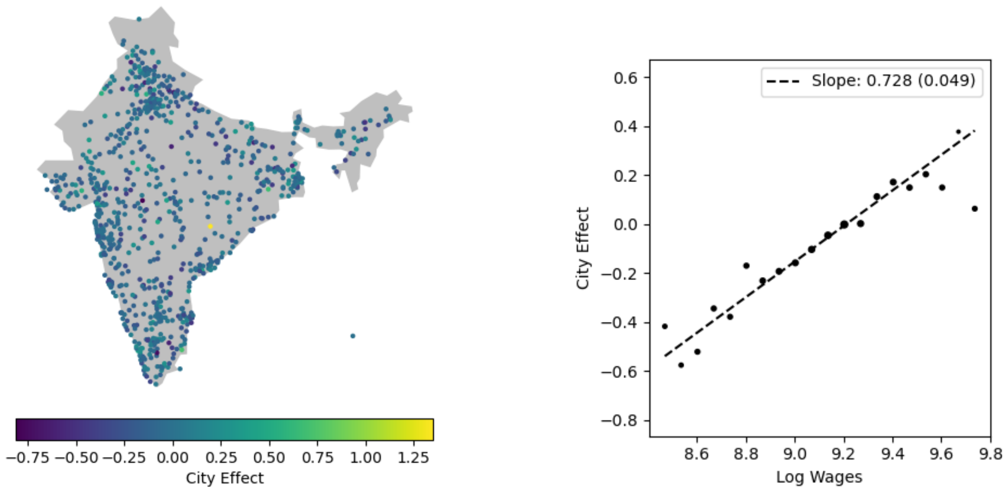

(a) Estimated City Effect

(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.

with cross-border moves substantially exceed those associated with internal moves. Moreover, the heterogeneity in slopes across countries suggests that the role of city effects in explaining wage differentials varies meaningfully by national context. Hence, the potential wage gains from internal migration may differ across countries, with larger gains where city effects account for a greater share of wage differences. We examine cross-country variation in city effects systematically below.

We can also aggregate city effects to the country level, weighting cities by population. This captures average differences in estimated place effects across countries. Appendix Figure A9 maps these country effects and shows transition matrices across countries. Individuals are more likely to migrate to countries with higher wages and higher country effects, consistent with migration moving workers toward higher-return locations.

5.2 The Contribution of City Effects to Wages

Table 1 decomposes cross-city wage variation worldwide 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). Columns (1)–(3) compare cities above and below the median, in the top and bottom quartiles, and in 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 Ω, the share of variance in mean city wages attributable to city effects, as defined in Equation 8. This statistic corresponds to the slopes shown in Figures 9, 10, and 11.

Table 1 shows that city effects account for 93% of global wage variation across cities. Applying

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					(0.8)

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 32.

the bound from Equation 32, we report a conservative value of 80% in parentheses in Column (5). Returning to the Bangalore-San Francisco example from the introduction, 93% of the wage difference is attributable to the difference in estimated city effects, echoing evidence on large cross-border place premia (Clemens et al., 2019).

For our sample of developed countries, city effects account for 45–51% of wage variation across cities (Table 2). Allowing for potential bias from wage imputation, Equation 32 implies a more conservative range of 39–44%. Differences in city effects account for 45–57% of the wage gap between below- and above-median cities in developed countries. In 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 wage gap.

For our sample of developing countries, city effects account for 48–73% of wage differences, with a conservative bound of 41–63% under potential imputation bias (Table 3). Overall, city effects explain a substantial share of wage variation across cities and across partitions of the wage distribution. Moreover, the share of wage variation explained by city effects is systematically higher in developing countries–a pattern also visible in Figure A10 and tested formally in Section 6.

Appendix Table A2 reports the full variance decomposition. The signs and magnitudes are broadly similar to those in Card et al. (2025), with some differences. For example, the covariance between worker and firm effects is positive in the US but negative in some countries, including India. These patterns may reflect differences in assortative matching and allocative efficiency, which we discuss further in Section 6.1. Additionally, we rerun our analysis with metro-area groupings to assess how sensitive our results are to a particular definition of a location or city. As reported in Appendix Table A3, we find similar results for our variance decompositions when we pool locations into metro areas and limit our sample to these locations before estimating location effects.

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					(0.44)

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					(0.43)

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					(0.39)

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 32.

5.3 Validating the AKM specification

As discussed in Section 4, interpreting our estimates as city effects requires exogenous mobility: moves across cities or firms must be uncorrelated with the error term ϵ_{it}. The event-study analysis provides evidence in support of this assumption. Here, we conduct additional tests using our estimated city effects, following Card et al. (2025).

Panel A of Figure 12 plots changes in earnings for cross-city movers against the origin-to-destination change in city effects. We measure earnings changes as earnings two years after the

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					(0.63)

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					(0.46)

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					(0.41)

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 bound for the explained share based on Equation 32.

move minus earnings two years before the move. The 45-degree line corresponds to earnings changes equal to the change in city effects. The scatter plot has a slope of 0.786, suggesting that movers experience, on average, 78.6% of the earnings change predicted by the change in city effects.

We decompose the earnings change for a worker who moves between periods $t-1$ and t:

$$y_{i,t} - y_{i,t-1} = \left(\Gamma_{c(i,t)} - \Gamma_{c(i,t-1)}\right) + \left(h_{f(i,t)} - h_{f(i,t-1)}\right) + \left(\epsilon_{i,t} - \epsilon_{i,t-1}\right) \ , \tag{15}$$

where Γ_c is the city effect, h_f is the hierarchy component of earnings based on the firm's position

Figure 12: Change in earnings components of city movers around move, against the change in city effect (within country movers)

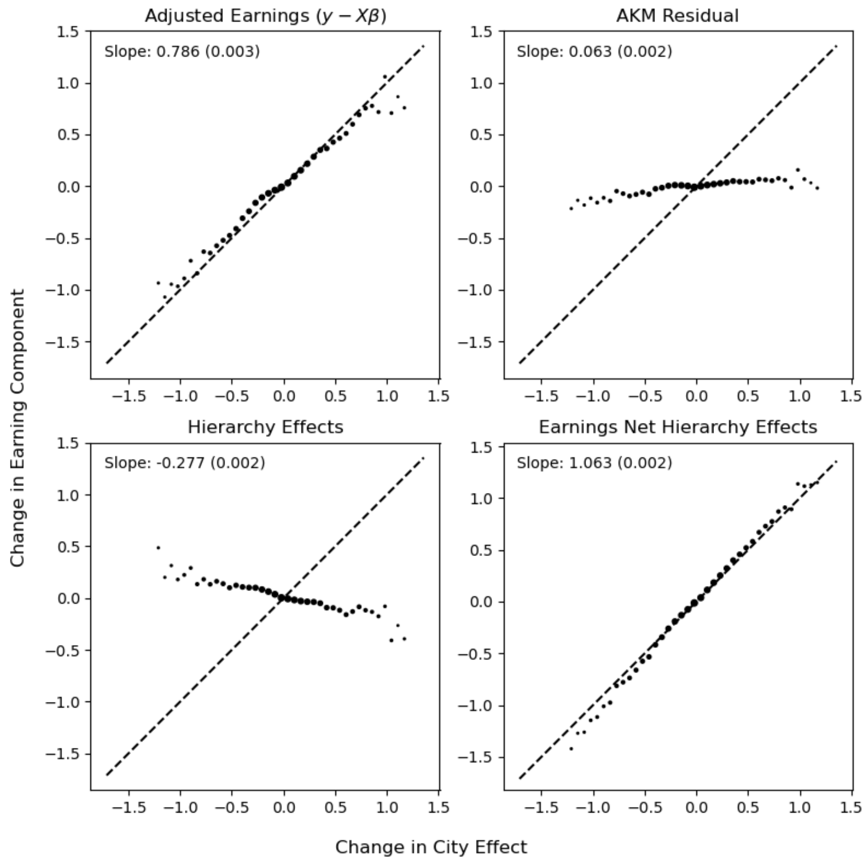

Notes: We look at changes in earnings two years before and after a move. X-axis plots the change in city effects for movers as destination city effect - origin city effect. City effects are constructed using Equation 4. We plot a 45' line as reference. The sample is restricted to movers who move once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

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, this would suggest a violation of exogenous mobility. Panel B tests this implication. The relationship is relatively flat, with a precisely estimated slope of 0.063. The small magnitude of this coefficient suggests that any departure from the assumption is quantitatively limited.

Panel C plots changes in the hierarchy component against changes in city effects. The negative correlation indicates that workers moving from low- to high-effect cities experience a decrease in the hierarchy component. These movers join firms lower on the local job ladder than the firms they leave. Workers moving in the opposite direction tend to join firms higher on the local job ladder.

Finally, Panel D shows that hierarchy-adjusted earnings changes track city effects closely, with a slope of 1.06. As in Card et al. (2025), this evidence suggests that violations of exogenous mobility are quantitatively small and that the AKM specification captures city effects well. We reach similar conclusions for the set of all movers and for city effects estimated using the connected set of international movers (Figure A11), except that hierarchy effects are quantitatively less important when comparing cities across countries. Table A4 shows that estimated city effects are similar whether we use all moves or restrict to the first two cities observed for each worker.

We conduct one additional exercise to assess whether housing costs affect estimated city effects. Using Zillow data for US cities, we deflate wages in each city by housing costs. We then re-estimate city effects using these deflated wages. Figure A7 shows an almost one-to-one relationship between city effects estimated with deflated and non-deflated wages.

6 What Explains the City Effects and Their Distribution

6.1 The Allocation of Workers Across Firms

City effects may partly reflect where high-premium firms locate (Carry et al., 2025) and how workers are allocated across those firms. We decompose estimated city effects to determine how much variation is driven by the presence of high-premium firms versus the allocation of workers across them. Building on Olley and Pakes (1996), we rewrite the city effect as:

$$\Gamma_j = \sum_{j(f)=j} s_f \gamma_f = \bar{\gamma}_j + \sum_{j(f)=j} (s_f - \bar{s}_j)(\gamma_f - \bar{\gamma}_j) \ , \tag{16}$$

$\bar{\gamma}_j$ is the average firm effect in city j and $\bar{s}_j$ is the average employment share in j. This decomposition separates the average firm effect in a city from the allocation of workers across firms.

If above-average firms employ a larger share of workers within a city, the allocation component raises the city effect. Panels (a)–(d) of Figure 13 plot the decomposition across the distribution of average city wages within countries, while Panels (e) and (f) split the allocation component by firm size to isolate the role of the largest firms.

In the US and UK, worker allocation across the firm-premium distribution is an important component of city effects: higher-effect cities employ more workers in high-premium firms, whereas lower-effect cities employ too few workers in those firms. By contrast, in India and Mexico, employment is less concentrated in high-premium firms, especially in lower-effect cities. In lower-effect cities, high-premium firms employ too few workers, while in higher-effect cities there is weak or no evidence that high-premium firms employ more workers.

This evidence is consistent with studies arguing that misallocation is a larger impediment to growth in developing countries. Here, we suggest that the allocation of workers *within* cities lowers measured city effects in developing countries. When we split the allocation component by firm size, we find that the largest, “superstar” firms drive much of the allocation component in the US, meaning that the largest firms are also among the highest-premium firms. In India, by contrast, the largest firms often have lower premia than the average firm within a city.

Figure 13: Decomposition of City Effects According to Employment Allocation

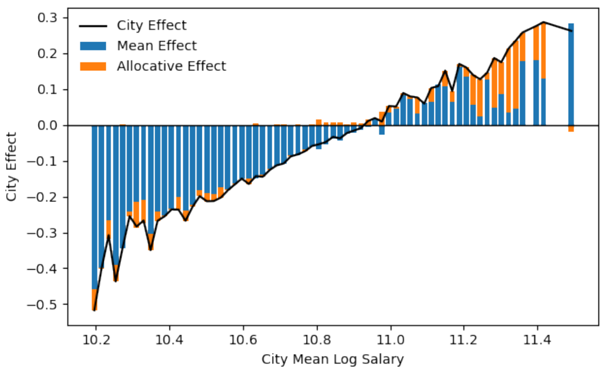

(a) USA Decomposition

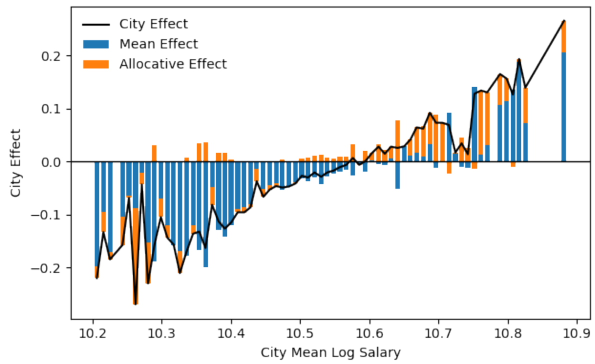

(b) United Kingdom Decomposition

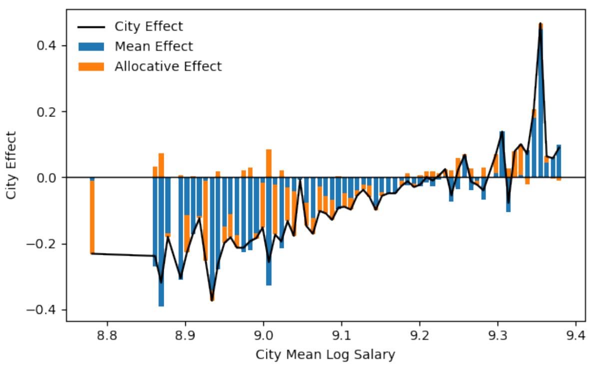

(c) India Decomposition

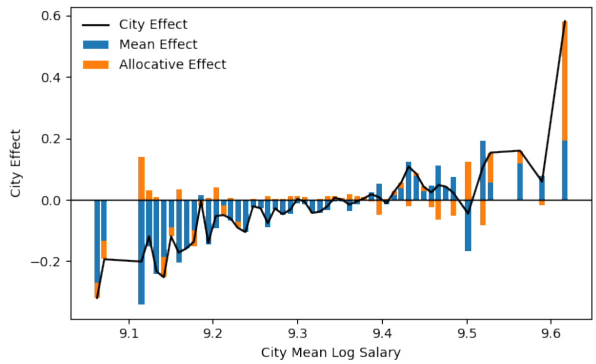

(d) Mexico Decomposition

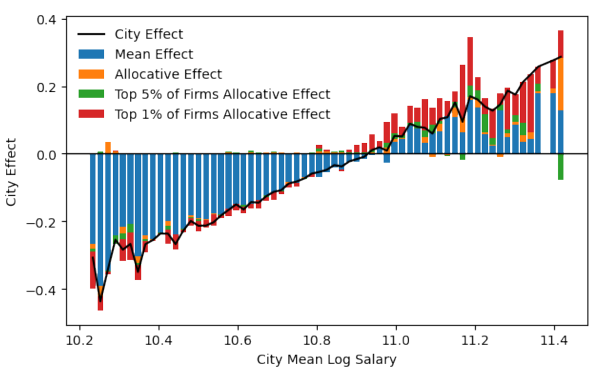

(e) USA Decomposition, with Superstar Firms

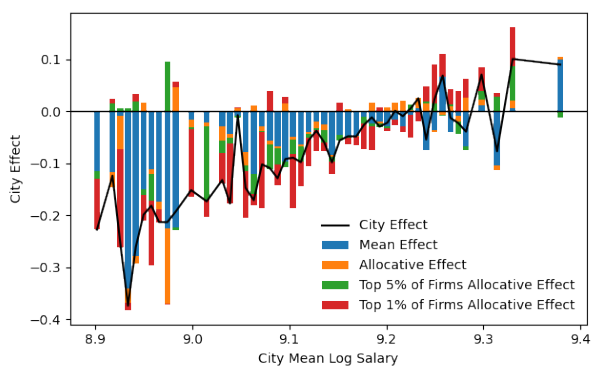

(f) India Decomposition, with Superstar Firms

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 (16). 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.

6.2 Explaining the City Effects

We next examine whether estimated city effects are associated with observable city characteristics. Using the full sample of LinkedIn users, we construct measures of industrial diversity, the share of senior job titles, average job tenure, and average remote-work suitability. We also obtain current city population by merging our cities with publicly available GeoNames data. We estimate the relationship between city effects and each characteristic, including region fixed effects to absorb broad regional heterogeneity.[19]

Figure 14a reports standardized coefficients, while the remaining panels present binned scatter plots for each descriptive (non-causal) relationship. Consistent with agglomeration, larger and denser cities have higher city effects, aligning with work emphasizing how density and scale facilitate knowledge spillovers and matching externalities (Rosenthal and Strange, 2004; Yi and Moretti, 2024). Cities with more remote-suitable jobs also have higher city effects.

We examine density directly by matching our cities to GHSL Functional Urban Areas. Figure 15a plots the relationship between density and city effects, while Figure 15b splits the relationship by development status. Denser cities have higher estimated city effects, and the relationship is stronger in developing countries. This stronger relationship is driven primarily by lower effects in less dense developing-country cities, rather than especially high effects in the densest cities. In Figure A12 we find a similar positive relationship between city effects and both city output and output per capita using global estimates of GDP from Rossi-Hansberg and Zhang (2026), further validating our estimates and drawing a stronger link to potential productivity measures.

We next examine skill-based agglomeration. Prior work argues that higher densities of skilled workers foster innovation (Moretti, 2004) and tenure helps explain differences in wage growth (Donovan et al., 2023). We capture skill intensity using job seniority, O*NET-based occupational classifications, the share of tech jobs, remote-work suitability, and job stability. Across dimensions, cities with more skilled, stable, and senior job structures have higher estimated city effects.

We then study industrial diversity in Figure 14b. Cities with lower industrial HHI exhibit larger city effects, consistent with complex, diversified economic structures that support sectoral linkages and network externalities. In contrast, concentrated industrial structures have lower city effects.

We also find substantial heterogeneity across country-level indicators. Appendix Figure A13 repeats these correlations after splitting countries by development status, urbanization, and college attainment. We find city size is more strongly associated with city effects in developing countries, as smaller cities in developing economies have especially low estimated effects.

We further explore the correlation with each industry in Appendix Figure A14. Computer science, management, legal, and other high-skill professional occupations are positively correlated with higher city effects. By contrast, retail, personal-care, and food-service occupations are negatively correlated with city effects. Our city effects are strongly associated with observable measures of urban economic structure and are robust to more granular fixed effects. These findings help characterize what distinguishes high-effect cities. High-effect cities tend to have more high-skill workers, more diverse industries, and more skilled, stable jobs.

[19] 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 14: Correlates with City Effects

(a) Correlates of City Effects

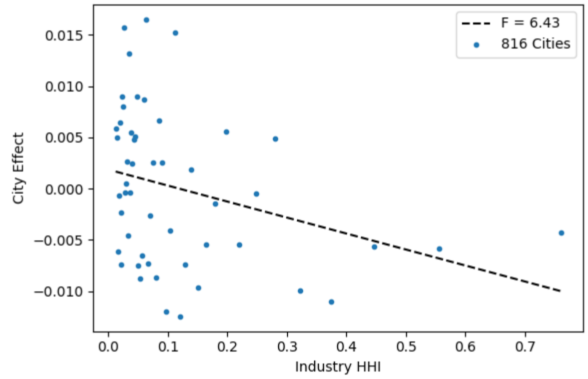

(b) Industrial Diversity

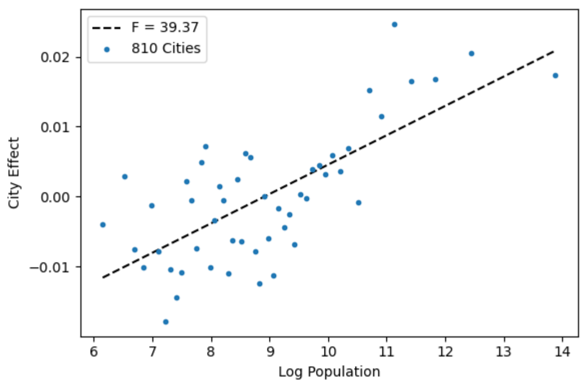

(c) Population

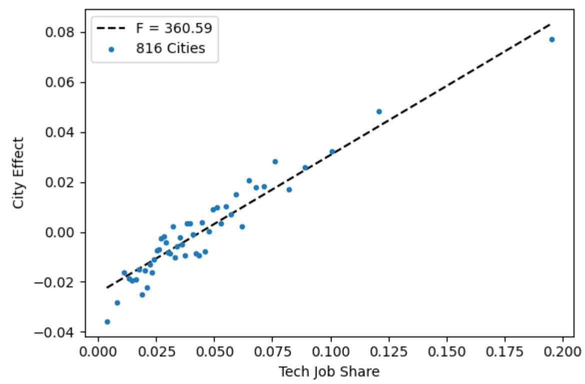

(d) Tech Job Share

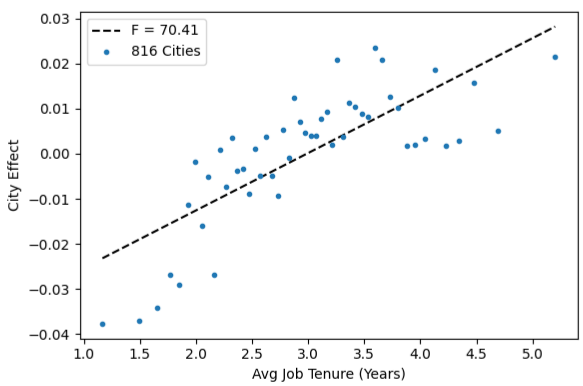

(e) Average Job Tenure

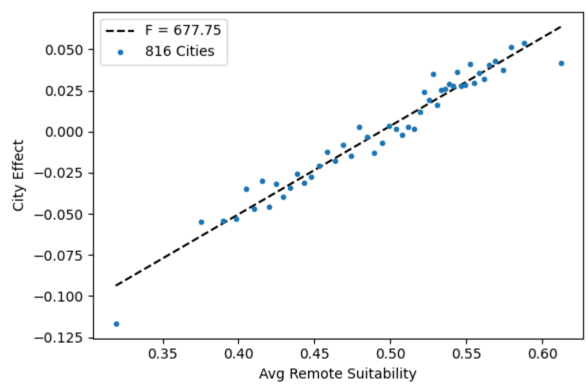

(f) Job Remote Suitability

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 15: Population Density and City Effects

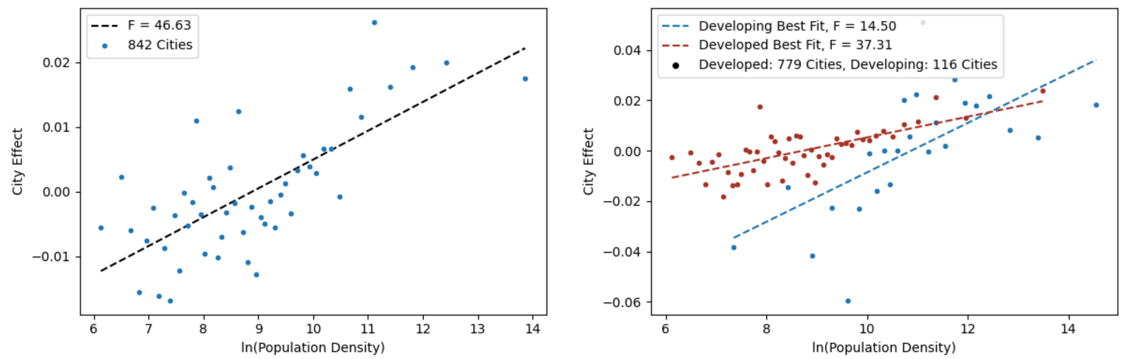

(a) Population Density (b) Population Density, by Development

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 population density of a city. Panel (b) plots the same relationship, but split by country development status.

Finally, we examine how city effects vary with proximity to high-effect cities. Figure 16a shows a strong negative relationship between a city's estimated effect and its distance to the nearest top-decile city in the country. This pattern is consistent with quantitative spatial models in which better market access raises local earnings or productivity. Importantly, the relationship is not driven solely by suburbs of major cities; it 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 highest-effect city, consistent with reduced competition and localized agglomeration gains.

Figure 16: Distance to Big Cities and City Effects

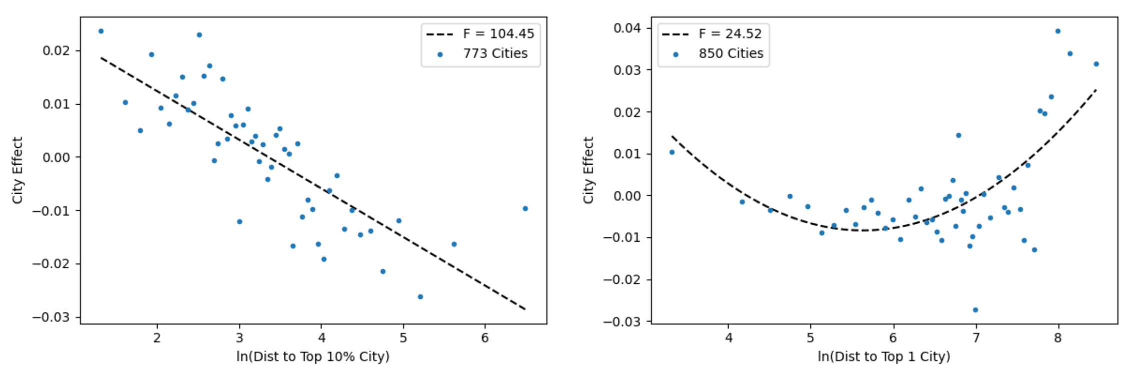

(a) Distance to Top 10% of City Effects (b) Distance to Highest City Effect

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.

6.3 Amenities and City Earnings Premia

Our estimates capture earnings differences, not welfare. We therefore examine how city earnings premia relate to urban (dis)amenities (Rosenthal-Kay, 2026). In a spatial equilibrium, attractive amenities may reduce the earnings premium required to draw workers to a city. However, these relationships do not necessarily identify compensating differentials, as city characteristics may be by-products of production or may themselves affect productivity. We therefore interpret the evidence as describing the equilibrium relationship between city characteristics and earnings premia.

Figure 17: Amenities and City Effects

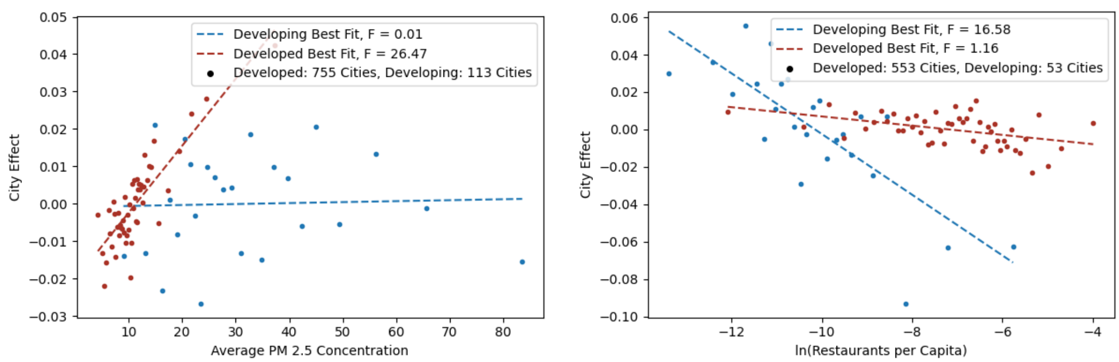

(a) PM 2.5 Concentration, By Development (b) Restaurants per Capita, By Development

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, split by country development status. Panel (b) plots the relationship between city effects and log restaurants per capita, split by country development status.

We begin with environmental externalities associated with economic activity. Figure 17a plots the relationship between annual PM 2.5 concentration and city effects, split by country development. Higher-effect cities tend to have more pollution, with a substantially steeper relationship in developed countries. This positive association is consistent with both an earnings premium that compensates workers for pollution, and pollution proxying for production intensity or congestion.[20]

We then turn to consumption amenities. Figure 17b shows that cities with more restaurants per capita have lower city earnings premia, with a steeper relationship in developing countries. This pattern suggests that amenity-rich cities may require lower earnings premia, although restaurants may also proxy for local economic structure.[21] Appendix Figure A15 presents analogous correlations between city effects and weather characteristics.

Overall, the results show that city earnings premia are systematically related to local amenities. Higher premia in more polluted cities and in cities with fewer restaurants are consistent with spatial compensation, although the correlations may also reflect production and urban structure.

[20]We find similar patterns when relating city effects to total earnings (population $\times$ wages), consistent with pollution in poorer countries being pervasive rather than concentrated in their highest-effect locations.

[21]The relationship remains negative when we control for region fixed effects and industry-specific employment.

7 Potential Aggregate Gains from Allocation Across Cities

7.1 Migration Toward High-Effect Cities

Our data allow us to measure how far individuals move, both within and across countries. We test whether longer-distance moves are associated with larger gains in city effects. If longer moves yield higher gains but few individuals undertake them, this may reflect strong preferences for remaining near one's origin or high migration costs that increase with distance.

The top panels of Figure 18 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 shown in Figure 16, these high-productivity cities are often farther away, consistent with the presence of "agglomeration shadows."

Figure 18: Move Distance Against City Effect Gain

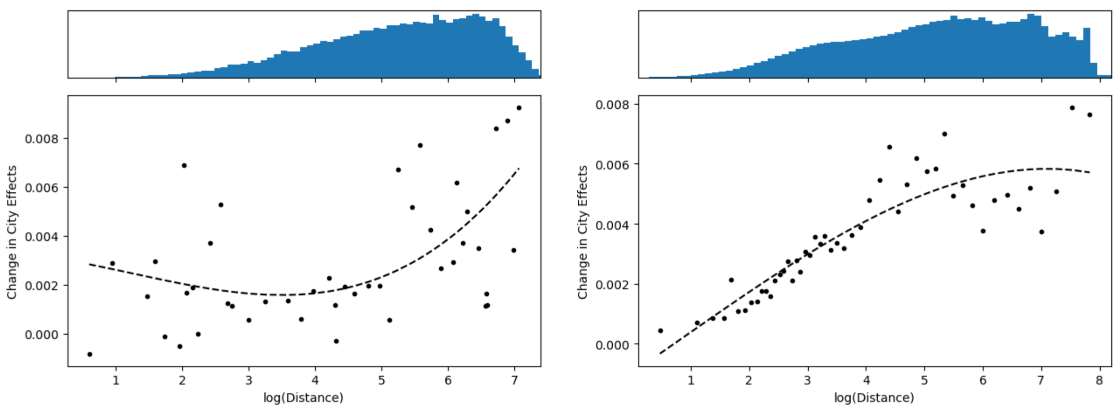

(a) Within Low Income Countries

(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 city-effect 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 larger gains in city effects. These patterns suggest that individuals who move long distances often capture larger returns, but geographic preferences and regional attachments may prevent many workers from reaching the highest-effect opportunities within their country.

There may be substantial potential gains from moving across cities. Such mobility may improve the allocation of labor and raise aggregate earnings when workers move from lower- to higher-effect cities. If wages are proportional to marginal products, these earnings gains may also translate into productivity gains. Firms may have monopsony power, however, and wages may be marked down.

If markdowns vary systematically by development status or city size, we should interpret 'potential gains' with caution. With these caveats in mind, we revisit the transition matrices, now analyzing migration across the *city-effect* distribution rather than across *wage* quintiles.

Figure 4 previously presented transition matrices for internal migration by city wage quintiles. Across most countries, the brightest cells appeared in the top row, indicating that movers disproportionately chose the highest-wage cities regardless of their origin quintile. This pattern held in both high-income countries such as the United States and lower-income settings such as India. If higher wages reflected larger city effects, these moves would generate meaningful allocative gains.

Figure 19: Transition Matrices for Internal Movers by City Effect Quintile

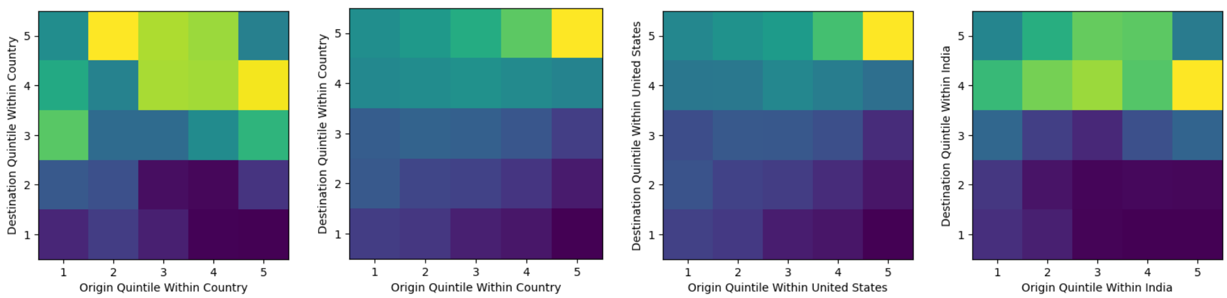

(a) Low Income Countries (b) High Income Countries (c) In the US (d) In India

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 19, 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 19a), 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 suggests that these destinations tend to be large, high-wage urban centers, while the highest-effect locations are often surrounding satellite towns. These contrasting patterns highlight that in low-income countries, particularly those with uneven spatial development, sizable potential earnings gains may remain from redirecting migration toward the highest-effect locations.

7.2 Heterogeneity of Variance of City Effects by Country

The global scope of our data allows us to assess whether potential returns to internal migration vary across countries. If the within-country distribution of city effects were the same across countries, workers moving from the 25th to the 75th percentile of their country's city-effect distribution would experience comparable wage gains regardless of country. When the city-effect distribution is narrower in one country than another, the returns to the same percentile move are lower in the country with less dispersion. This highlights the role of cross-country heterogeneity in city-effect dispersion in shaping the potential wage benefits of internal migration and the returns to facilitating migration (Bryan and Morten, 2019).

Figure 20: Distribution of City Effects

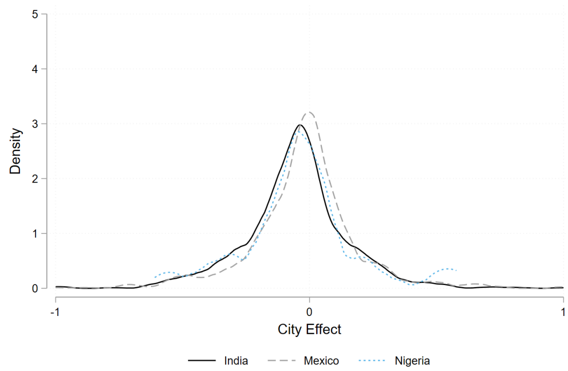

(a) City Effects Distribution, LMICs

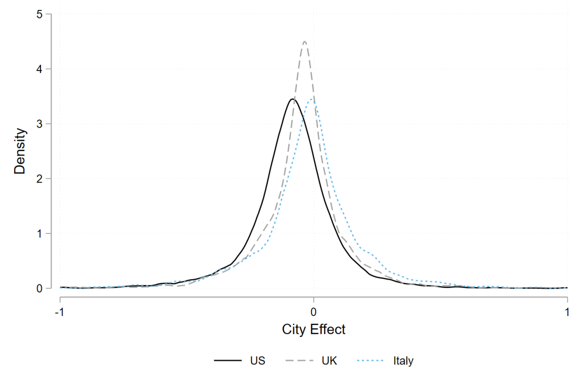

(b) City Effects Distribution, Developed Countries

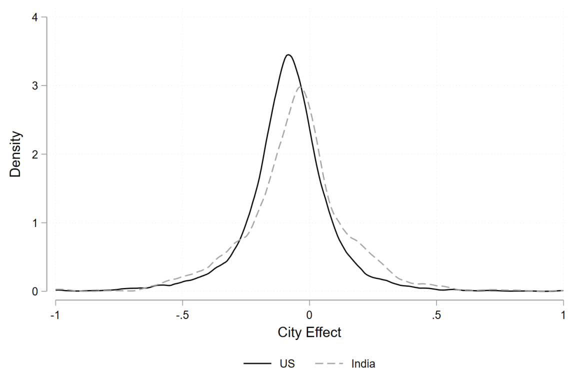

(c) City Effects Distribution, USA and India

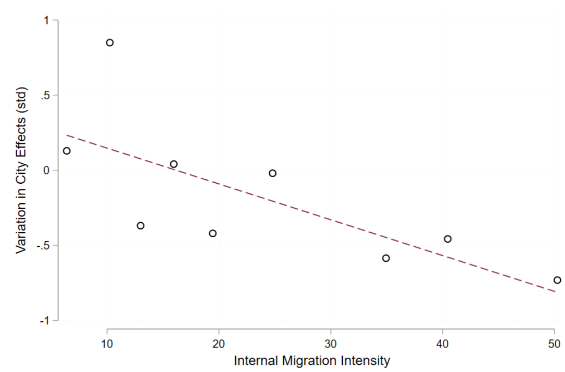

(d) Migration Intensity and Variation in City Effects

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.

To highlight cross-country differences in city-effect distributions, Panels (a)–(c) of Figure 20 plot the distributions for the US and India, and for samples of developed and developing countries. Two patterns emerge. First, city-effect distributions differ even within developed and developing-country groups. Second, city-effect distributions are generally wider in developing countries than developed countries. This wider dispersion suggests larger potential gains from reallocating workers across cities, especially if wages are correlated with productivity (Donovan et al., 2023).

A natural next step is to examine how migration relates to the dispersion of city effects. If migration barriers are high, workers are less able to sort efficiently across space. City-effect dispersion may therefore be larger when individual sorting across cities is more limited. Using harmonized cross-country measures of migration intensity from Bell et al. (2015), we estimate the relationship between internal migration and within-country city-effect dispersion.[22] Figure 20d shows a strong negative relationship between a country's internal migration rate and the variance of city effects.

These patterns imply that the potential returns to internal migration vary substantially across countries. They are largest where city effects are most dispersed, especially in developing economies, and smaller where the city-effect distribution is compressed. We next examine which country-level factors are associated with this dispersion.

7.3 The Importance of City Effects and Firm Frictions

Figure 21: Determinants of the Importance of City Effects

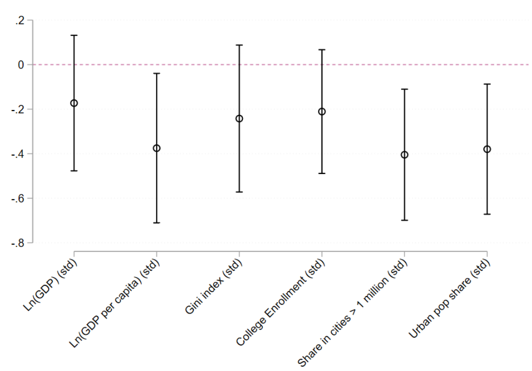

(a) Country Aggregates and Median Share

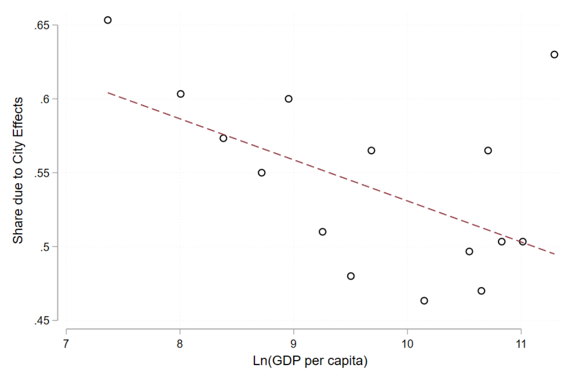

(b) GDP Per Capita and Median Share

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.

We next examine which country-level characteristics are associated with the importance of city effects. Using World Bank Development Indicators averaged over the previous 10 years, Figure 21a relates income, size, human capital, and urbanization to the share of wage differences explained by city effects, controlling for country coverage rates.

Wealthier, more educated, and more urbanized countries exhibit a smaller share of wage variation explained by city effects. Because a higher city-effect share implies larger potential gains from reallocating workers toward high-effect cities, these patterns point to larger returns to internal migration in low- and middle-income countries.

[22] Bell et al. (2015) cover 96 countries. In practice, we limit the sample to countries with coverage rates above 0.1, more than 50 identified cities, and over 20,000 unique users.

Figure 21b shows a strong negative relationship between GDP per capita and the share of wage differences explained by city effects, especially among low- and middle-income countries. The relationship weakens among high-income economies, possibly because housing supply constraints and land-use regulations become more important (Hsieh and Moretti, 2019; Harari and Wong, 2025b). Interpreting city effects as earnings-relevant place effects, the largest potential income gains from reducing spatial barriers appear in poorer countries, though mobility frictions may remain important even in advanced economies.

Firm-level frictions are also associated with a smaller role for city effects. Using World Bank B-READY data, Figure 22a shows that city effects explain less wage variation in countries with worse overall business environments, while Figure 22b reports results for specific components. Limitations on firm entry and utility quality are the strongest contributors: higher entry costs for domestic and foreign firms, and less reliable electricity supply, are both associated with a smaller share of wage variation explained by city effects. The electricity result is consistent with Fried and Lagakos (2023), who estimate large general-equilibrium losses from pervasive power outages.

Figure 22: Determinants of the Importance of City Effects

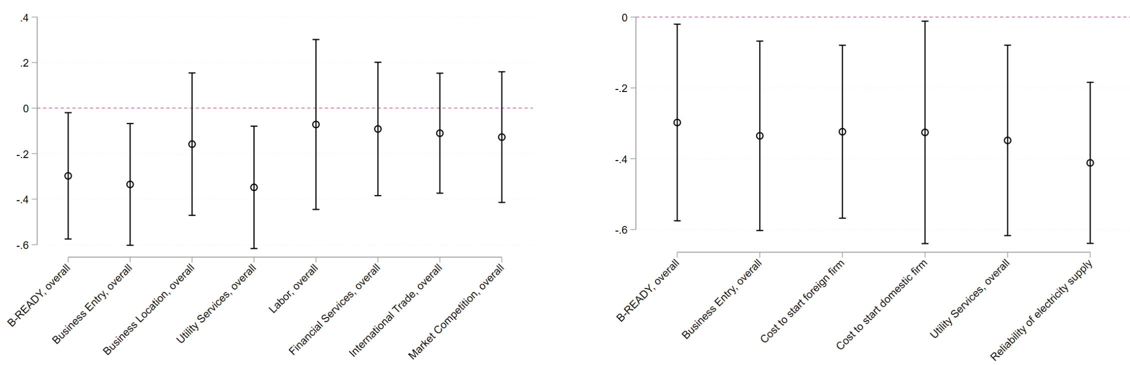

(a) B-READY Category Scores

(b) Firm Entry and Electricity

Panels (a) and (b) plot 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.

8 Quantifying Gains from Reallocation Within and Across Cities

We now use the estimated city effects to quantify potential wage gains from reallocating workers across and within cities. We consider three partial-equilibrium counterfactuals: reallocating workers across cities to match the US relationship between city effects and city size, improving within-city firm allocation to match the US, and doing both. These exercises require an earnings-productivity link and abstract from urban costs that could attenuate the gains (Rosenthal-Kay, 2026).

For the across-city exercise, we rank cities by city effects within each country and compute the population CDF across city-effect ranks.[23] We then reallocate each country's population to match the US CDF. Figure 23a compares India's observed and counterfactual distributions with the US distribution. Several large Indian cities fall just above the 65th percentile of the city-effect distribution rather than near the top, indicating scope for gains from reallocation.

To compute wage gains, we assign all workers in a city the city's average worker component before reallocation. This lets us randomly assign migrants from origin to destination cities and recalculate expected earnings after reallocation, while holding both the national average worker component and estimated city effects fixed. Any wage gains therefore come only from reallocating workers toward higher-effect cities. Even under this conservative setup, average annual wages rise by 2.3% in India, about USD 200 per worker, and by 4.0% and 5.9% in Mexico and Nigeria, respectively. The gains are generally smaller in developed countries, but not dramatically so: the average gain across Germany and the UK is 3.8%.

Because the full reallocation requires moving a large share of the population, we also consider a more modest counterfactual. We move workers from below- to above-median city-effect locations until the above-median population share matches the US, uniformly scaling each rank-based bin so that the within-half shape of the distribution stays fixed. In India, this moves only 8% of the population but raises average wages by 1.36%, achieving nearly two-thirds of the gains from the full reallocation. The contrast with developed countries is stark: in the UK 0.7% and in Germany 0.9% of workers move, and corresponding wage gains are only about 0.1% in both countries.

The within-city exercise draws on the decomposition in Section 6.1, which shows that workers in countries such as the US are better allocated to high-premium firms *within* locations. Rather than moving workers across cities, we reallocate the same number of workers across firms within cities. We group US cities by city-effect rank, calculate the average allocative effect for each rank bin, and apply those values to the mean effect, $\bar{\gamma}_j$, for cities in other countries.[24]

Figure 23b plots observed and counterfactual city effects for India against the US. Counterfactual city effects are almost uniformly higher than observed values in India, indicating that average allocative effects are higher in US cities. The increases are largest at the bottom and top of India's city-effect distribution, highlighting the potential for within-city reallocation to improve allocation within a country. Average wages in India rise by 2.6%. The gains are smaller in Mexico and Nigeria, at 1.6% and 2.9%, respectively. Consistent with diminishing scope for firm-level allocative improvements over the development process, gains are modest in high-income countries such as the UK (0.9%) and Germany (1.0%).

The final exercise combines the two margins: we first modify city effects using the US allocative effects and then reallocate population according to the modified city effects. Although the two

[23] 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.

[24] In practice, this adjustment yields smaller city effects than the mean effects for lower-effect cities and larger effects for higher-effect cities. Whether these counterfactual city effects exceed or fall below the estimated values depends on the relative magnitude of allocative effects across countries.

Figure 23: Quantifying the Gains from Re-allocation

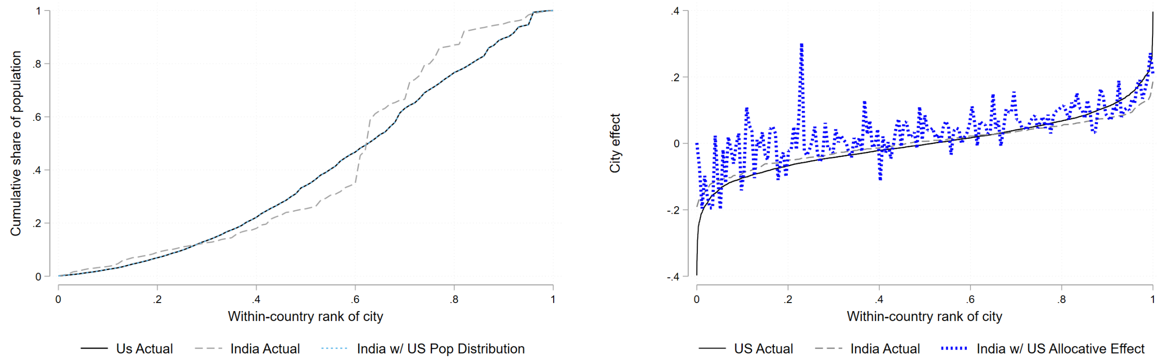

(a) Population Reallocation

(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.

mechanisms could offset one another if improvements in one margin reduce the gains from the other, we find little evidence of substitution. For India, the combined exercise raises average wages by 4.3%, with roughly 90% of the gains explained by the sum of the two exercises considered separately. The corresponding gains are slightly higher in Mexico (5.2%) and Nigeria (8.2%). These results suggest that policies reducing spatial and firm-level frictions need not compete, and that addressing both dimensions simultaneously can yield substantial aggregate income gains.

9 Conclusion

Cities are central to how economies organize production, allocate workers, and generate opportunity. Yet most evidence on the value of cities comes from a small number of high-income countries. We bring these questions to a global setting by estimating city effects from detailed job histories for 513 million workers in 220,000 cities across 191 countries. Event-study evidence from moves shows that earnings, seniority, and occupation- and industry-based wage measures change systematically when workers move across cities. We use these moves to separate city effects from ability-based sorting. Within countries, city effects explain 45–73% of cross-city wage differences, showing place remains a central determinant of earnings even after accounting for selection.

The estimated city effects allow us to revisit fundamental questions in urban economics across the development spectrum. High-effect cities are larger, denser, more industrially diverse, and more intensive in high-skill work. These relationships are consistent with agglomeration forces, but they vary sharply across countries. City-effect distributions are wider in poorer countries, and the importance of city effects declines with GDP per capita, education, and urbanization. This pattern is consistent with larger returns to internal migration in low- and middle-income countries.

City effects also reflect how workers are allocated across firms within cities. In richer countries, workers are more concentrated in high-premium firms, raising average city effects relative to poorer countries. Countries with weaker business environments, higher entry costs, and less reliable electricity also have a smaller share of wage variation explained by city effects. These margins matter quantitatively. In partial-equilibrium accounting exercises, reallocating workers toward higher-effect cities generates sizable aggregate wage gains, especially in developing countries. Applying the US within-city allocation pattern to other countries yields additional gains, and combining the across-city and within-city exercises produces roughly additive effects. Taken together, global city effects provide a useful empirical object for studying how urban structure, worker allocation, and firm-level frictions jointly shape economic opportunity around the world.

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.

——— (2026) "How Do Neighborhoods and Firms Affect Intergenerational Mobility?"Technical report, Working Paper, https://m0ises.github.io/mywebsite/CRY2026-IGE.pdf.

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.

——— (2013) "Why Do Programmers Earn More in Houston Than Hyderabad? Evidence from Randomized Processing of US Visas," *American Economic Review*, 103 (3), 198–202, 10.1257/aer.103.3.198.

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.

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.

Donovan, Kevin, Will Jianyu Lu, and Todd Schoellman (2023) "Labor Market Dynamics and Development," *The Quarterly Journal of Economics*, 138 (4), 2287–2325, 10.1093/qje/qjad019.

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

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.

Fried, Stephie and David Lagakos (2023) "Electricity and Firm Productivity: A General-Equilibrium Approach," *American Economic Journal: Macroeconomics*, 15 (4).

Gechter, Michael and Nick Tsivanidis (2023) "Spatial Spillovers from High-Rise Developments: Evidence from the Mumbai Mills," Working paper, revise and resubmit at Econometrica.

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.

Harris, Ian, Timothy Osborn, Phil Jones, and David Lister (2020) "Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset," *Scientific Data*, 7, 109.

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.

Morten, Melanie and Jaqueline Oliveira (2024) "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," *American Economic Journal: Applied Economics*, 16 (2), 389–421, 10.1257/app.20180487.

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.

Rosenthal-Kay, Jordan (2026) "Urban Costs Around the World," *Working Paper*.

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.

Rossi-Hansberg, Esteban and Jialing Zhang (2026) "Local GDP Estimates Around the World," *Journal of Urban Economics*, 154.

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

Shen, Siyuan, Aaron van Donkelaar, Nathan Jacobs, Chenguang Wang, and Randall Martin (2024) "Enhancing Global Estimation of Fine Particulate Matter Concentrations by Including Geophysical a Priori Information in Deep Learning," *ACS EST Air*, 1 (5), 332–345.

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.

Yi, Moises and Enrico Moretti (2024) "Size Matters: Matching Externalities and the Advantages of Large Labor Markets," NBER Working Paper 32250, National Bureau of Economic Research, 10.3386/w32250.

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

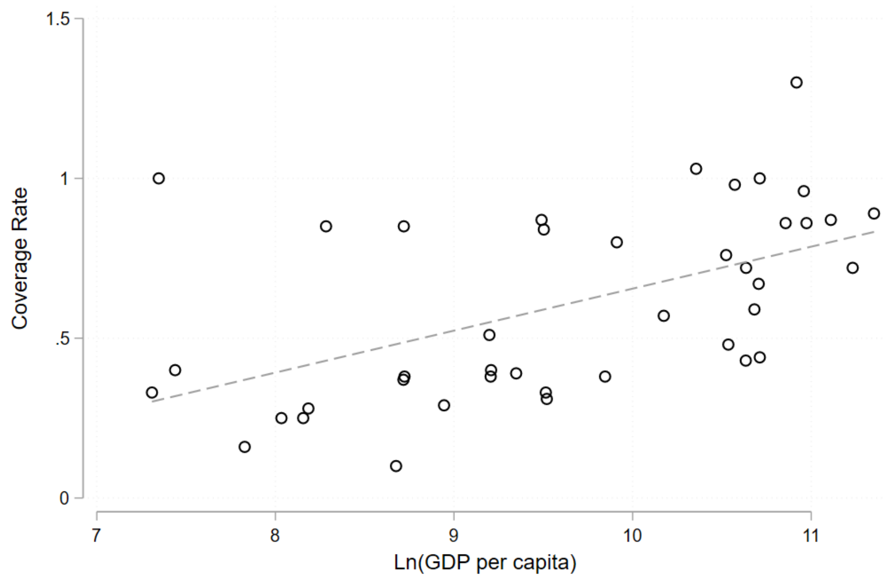

Figure A1: Coverage Rate, by GDP per Capita

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: Average City Wages in the UK and Italy

(a) UK

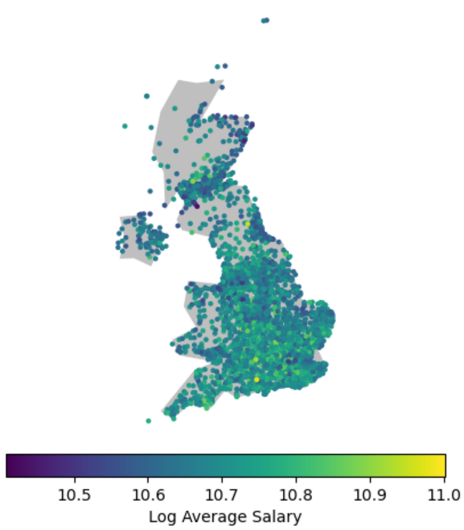

(b) Italy

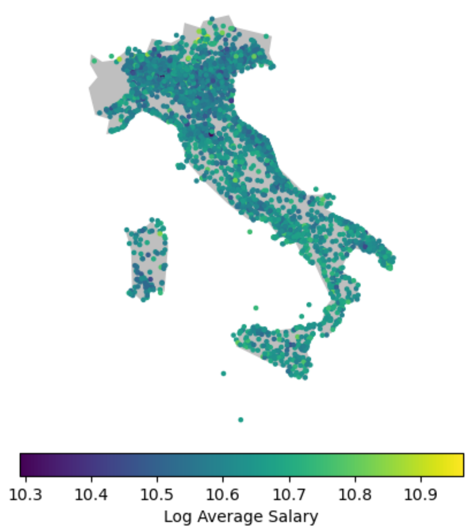

Notes: The maps plot average log(salaries) in each city. Legend scale varies by country. Sample restricted to cities with more than 100 LinkedIn users.

Figure A3: Event Studies by Outcome for Different Panels: Within-Country Movers and All Movers

Panel A: All Movers

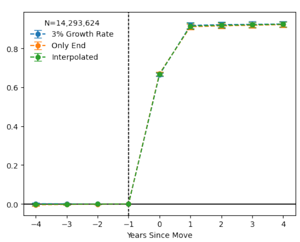

(a) Log Salary

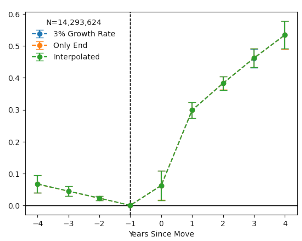

(b) Seniority

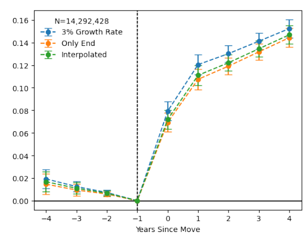

(c) ONET Wage

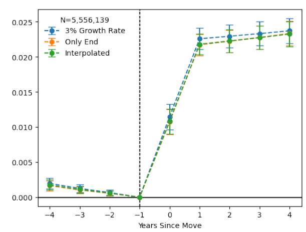

(d) NAICS Wage

Panel B: Within-Country Movers

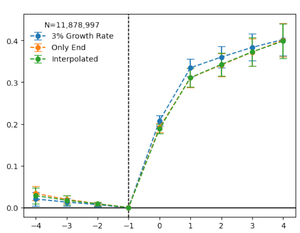

(e) Log Salary

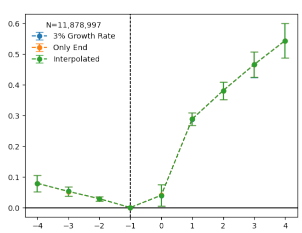

(f) Seniority

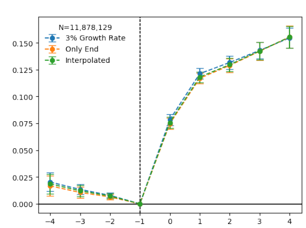

(g) ONET Wage

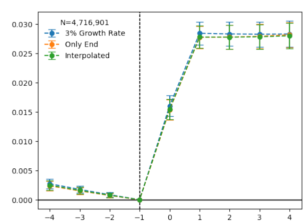

(h) NAICS Wage

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 A4: Event Study for Internal Movers for Different Panels, Developed Economies

Panel A: United States

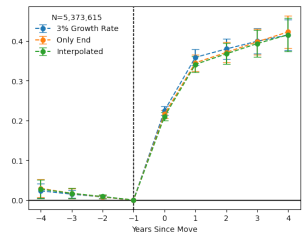

(a) Log Salary

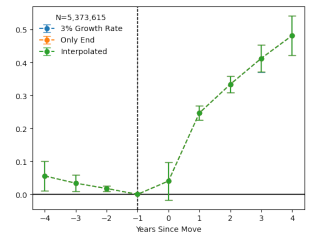

(b) Seniority

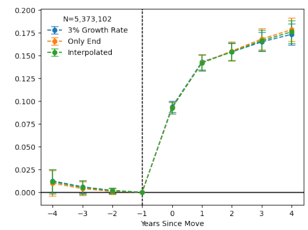

(c) ONET Wage

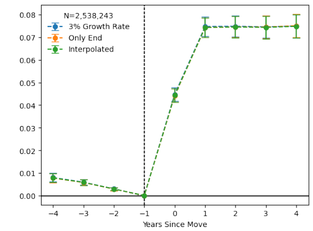

(d) NAICS Wage

Panel B: United Kingdom

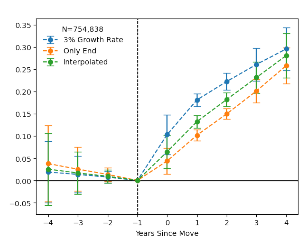

(e) Log Salary

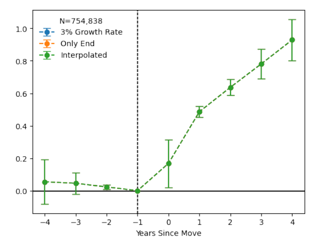

(f) Seniority

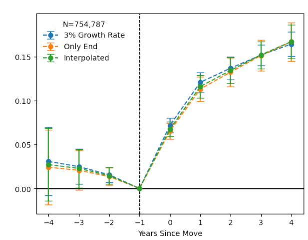

(g) ONET Wage

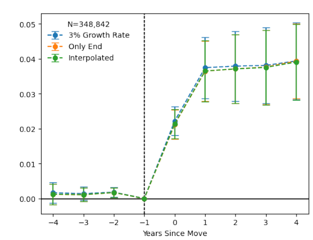

(h) NAICS Wage

Panel C: Italy

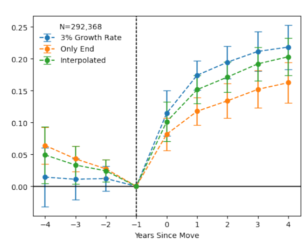

(i) Log Salary

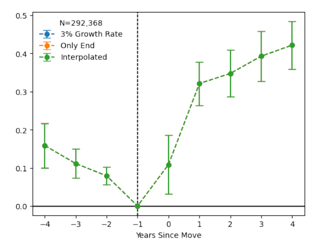

(j) Seniority

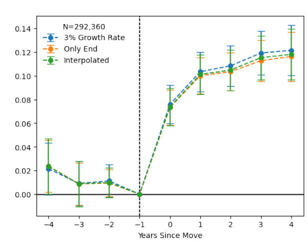

(k) ONET Wage

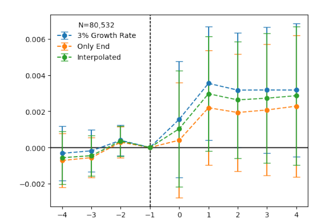

(l) NAICS Wage

Panel D: Germany

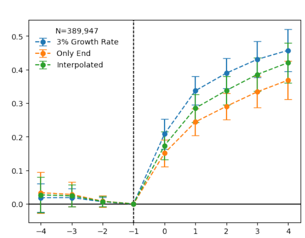

(m) Log Salary

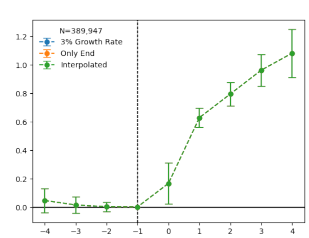

(n) Seniority

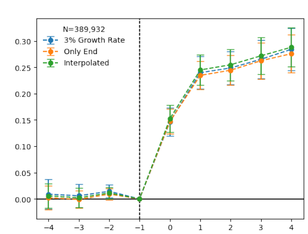

(o) ONET Wage

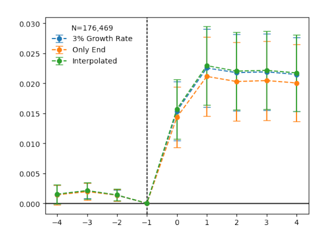

(p) NAICS Wage

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 A5: Event Study for Internal Movers for Different Panels, Developing Economies

Panel A: India

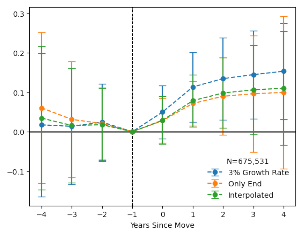

(a) Log Salary

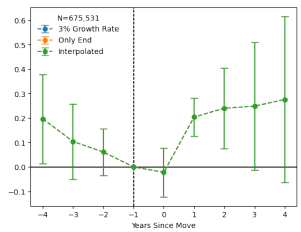

(b) Seniority

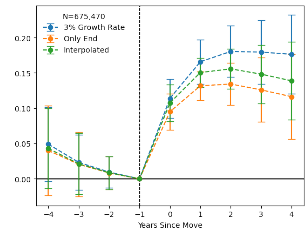

(c) ONET Wage

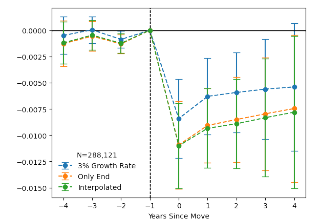

(d) NAICS Wage

Panel B: Mexico

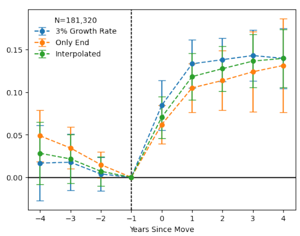

(e) Log Salary

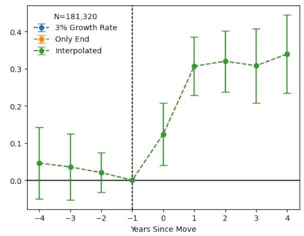

(f) Seniority

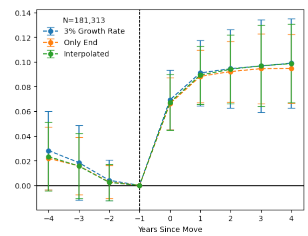

(g) ONET Wage

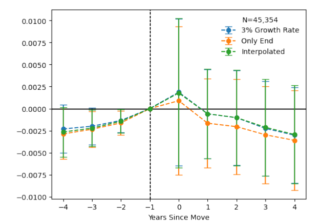

(h) NAICS Wage

Panel C: Nigeria

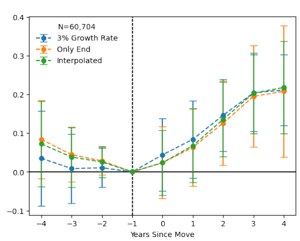

(i) Log Salary

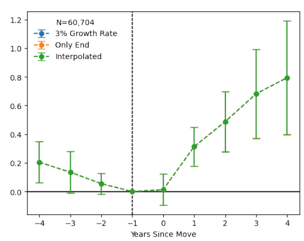

(j) Seniority

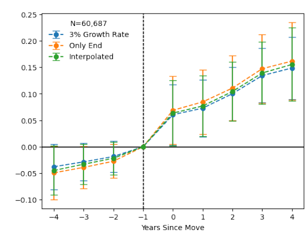

(k) ONET Wage

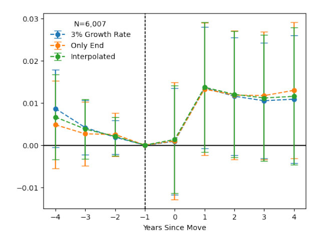

(l) NAICS Wage

Panel D: Chile

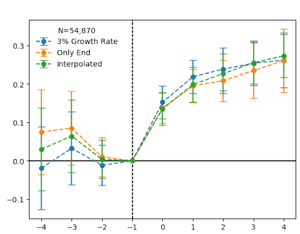

(m) Log Salary

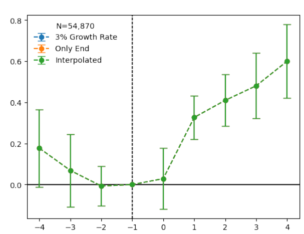

(n) Seniority

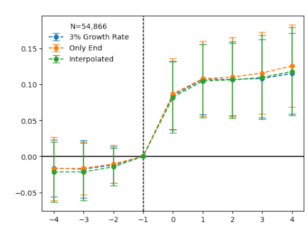

(o) ONET Wage

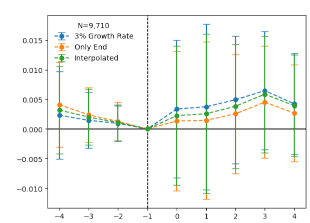

(p) NAICS Wage

Analogous to Figure 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 who worked in the year prior to moving ($t = -1$). Standard errors clustered at the user level.

Figure A6: City Effects Distribution for Different Panels

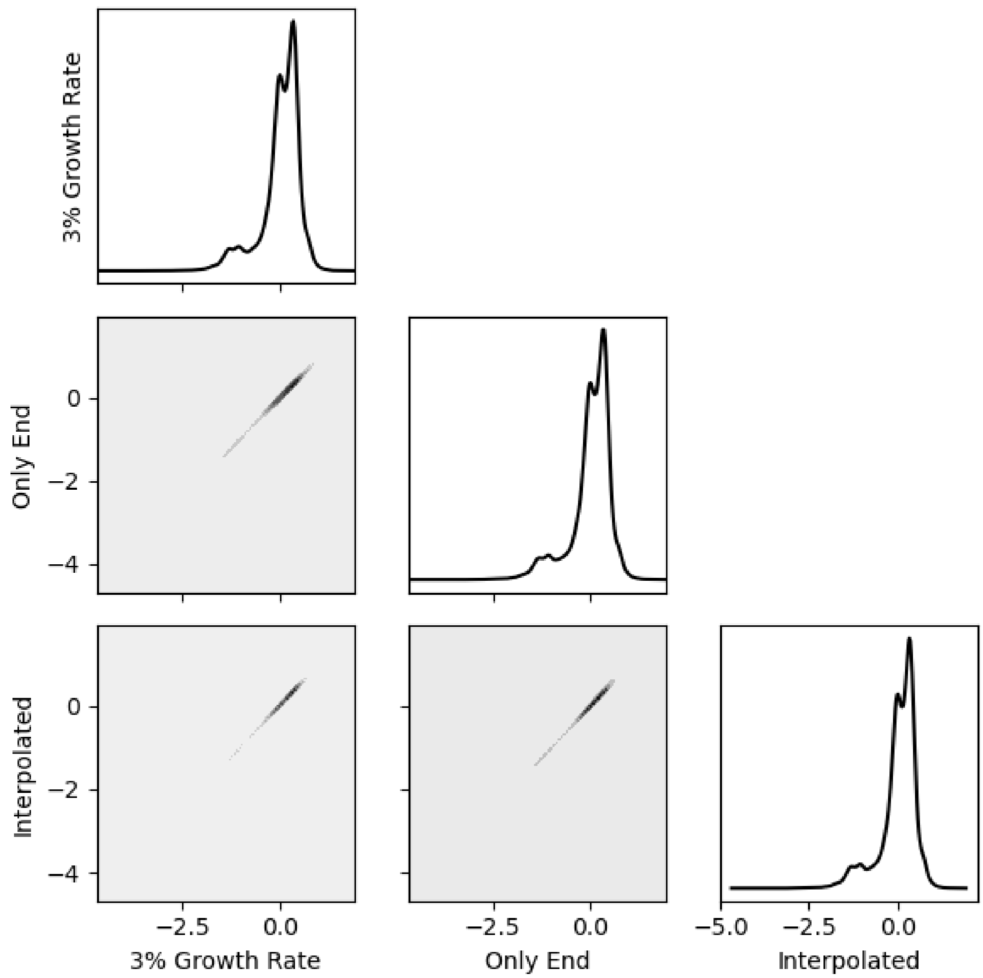

Distribution of city effects for different panels. The diagonal shows the univariate distribution of city effects. Below the diagonal shows the bivariate distribution of city effects with different panels. The “3% Growth Rate” panel, which is used in the paper, applies a 3% growth rate to user salaries 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 A7: City Effects Deflated by House Prices

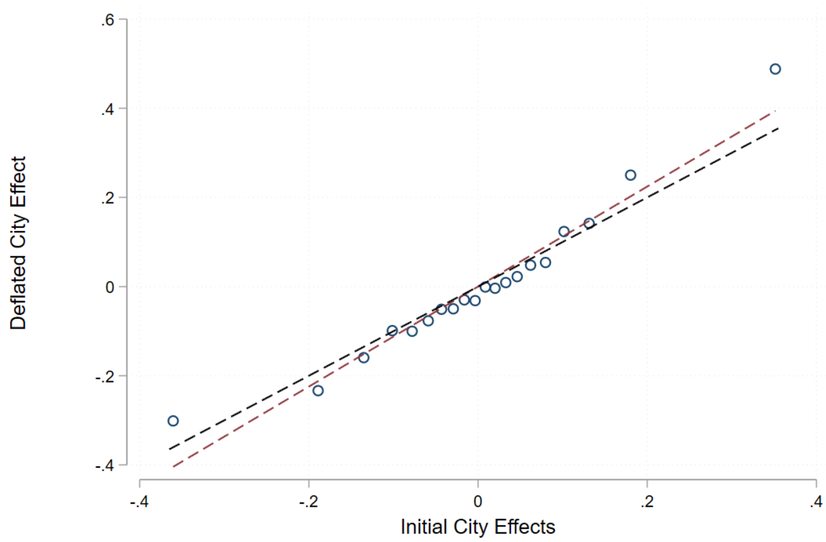

Notes: We plot the relationship between our original city effects we estimate against the city effects we obtain after deflating wages by house prices. This sample only includes US cities, and both city effects have been recentered and residualized by state FE before plotting. A 45° degree line is added in black.

Figure A8: Changes in Wages by Difference in City Wages, Over Time

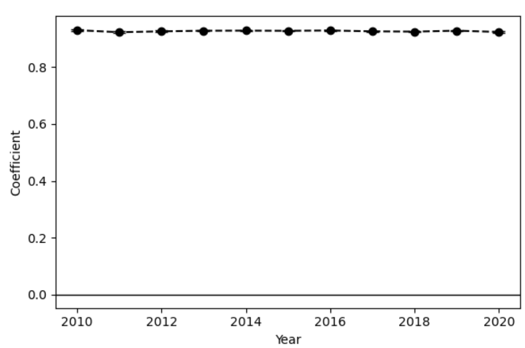

(a) International Moves

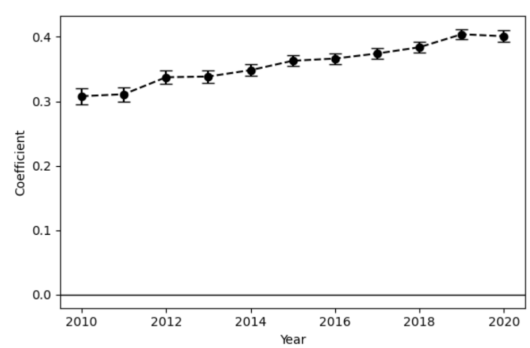

(b) Internal Moves

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 A9: Country effects

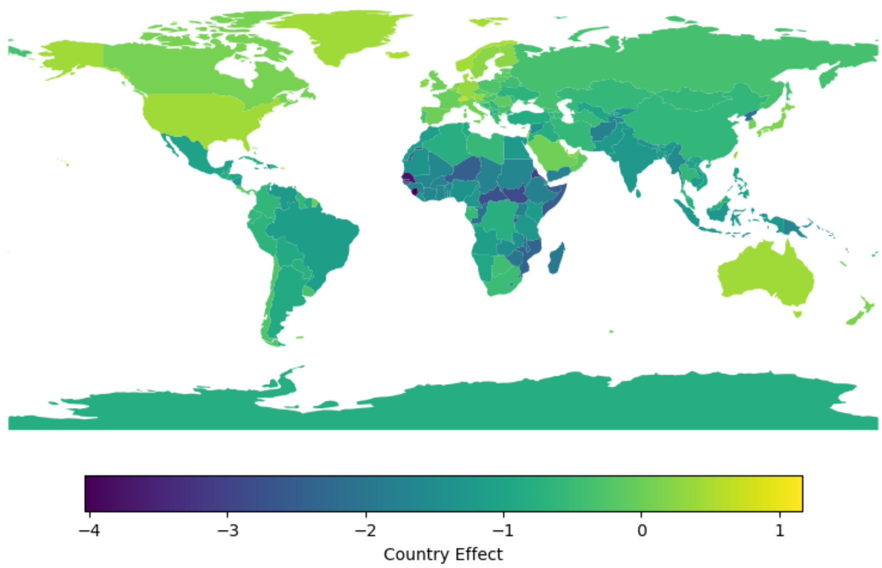

(a) Map of country effects

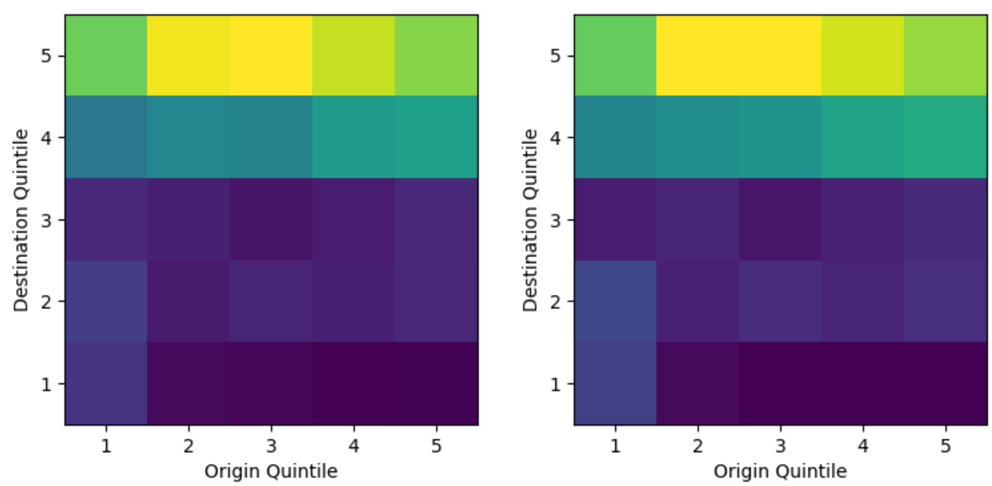

(b) Country salary transition matrix

(c) Country effect transition matrix

Country effects are calculated as the average of city effects within the country weighted by the user-years worked in that city (the same procedure used to create city effects to firm effects). Panel a maps the country effects. Panels (b) and (c) show transition matrices for the origin and destination salary and country effect quintiles of cross-country moves.

Figure A10: City Effects, Other Countries

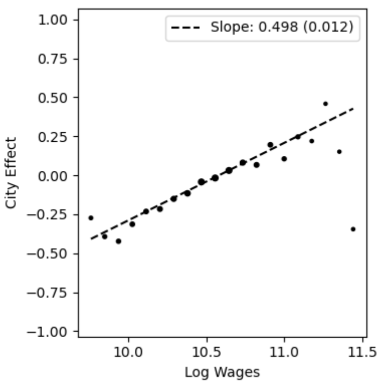

(a) UK

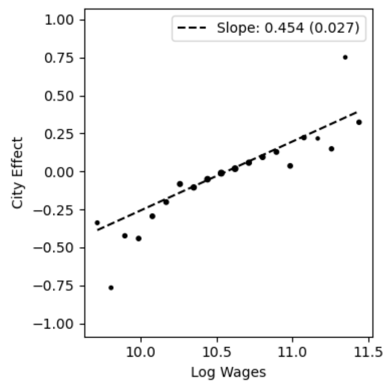

(b) Italy

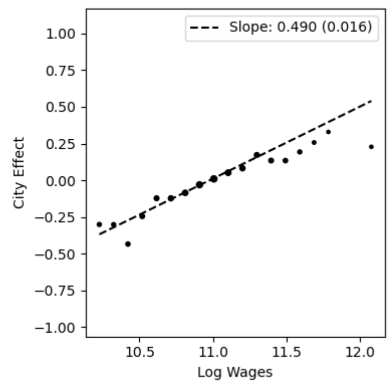

(c) Germany

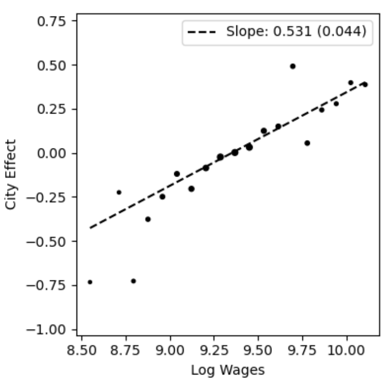

(d) Mexico

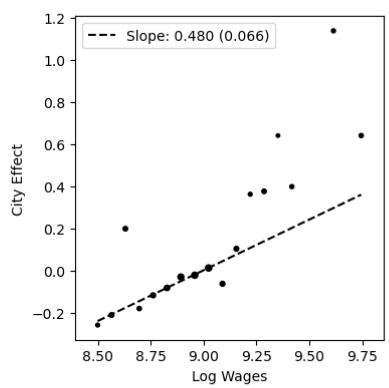

(e) Nigeria

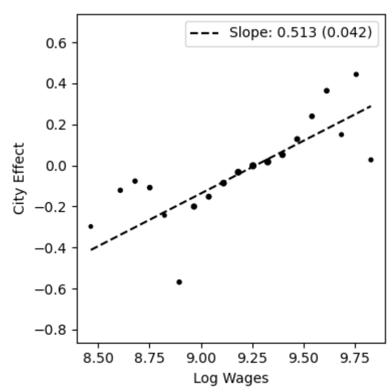

(f) Philippines

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 A11: Change in earnings components of city movers around move, against the change in city effect (all movers)

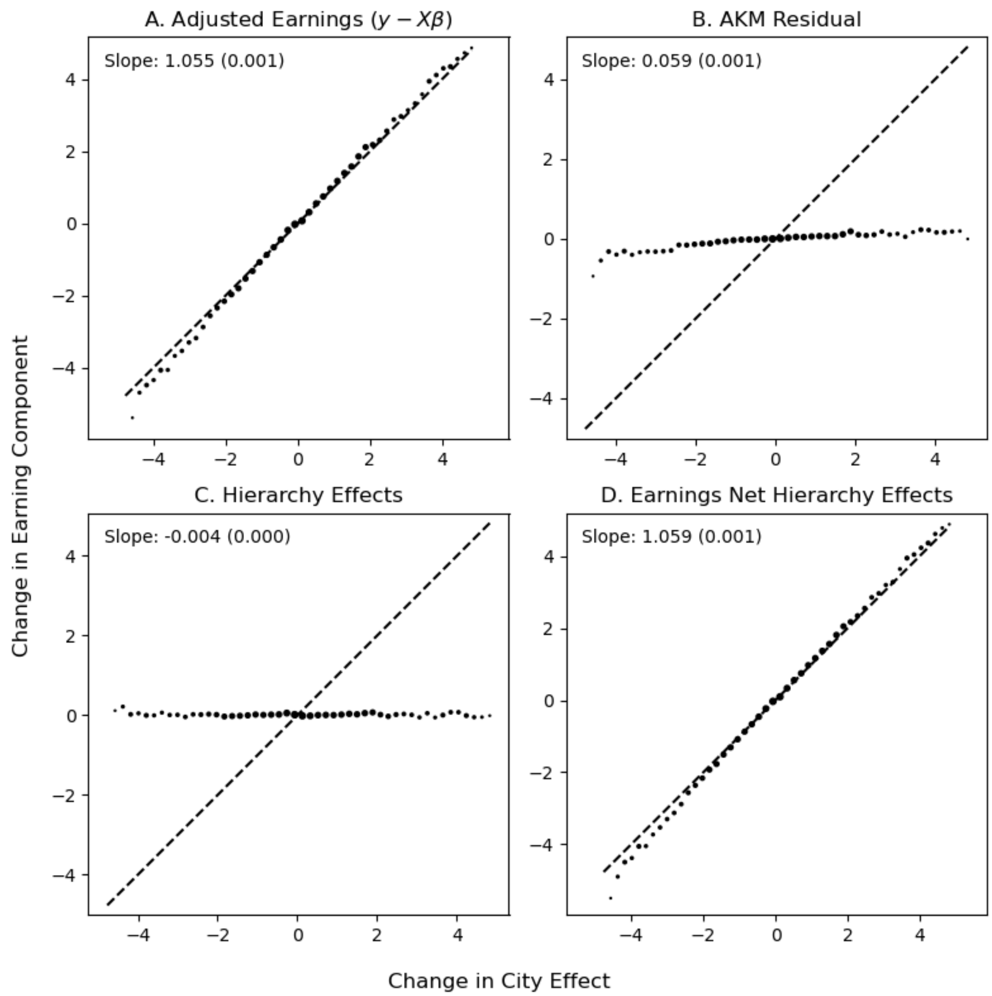

Notes: We look at changes in earnings two years before and after a move. X-axis plots the change in city effects for movers as destination city effect - origin city effect. City effects are constructed using equation 4. We plot a 45' line for reference. The sample is restricted to movers who move once.

Figure A12: GDP and City Effects

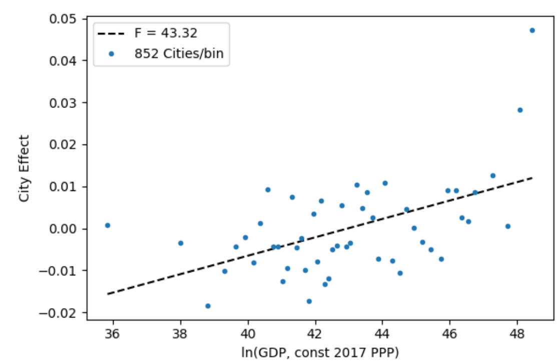

(a) GDP

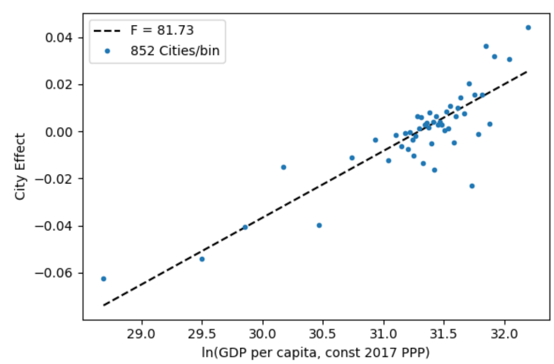

(b) GDP per capita

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 gross domestic product of all pixels of a city. Panel (b) plots the relationship between city effects and the gross domestic product per capita of all pixels of a city. Data from Rossi-Hansberg and Zhang (2026) and city boundaries are defined in the same way as in Figure 15.

Figure A13: Correlates with City Effects By Country Characteristics

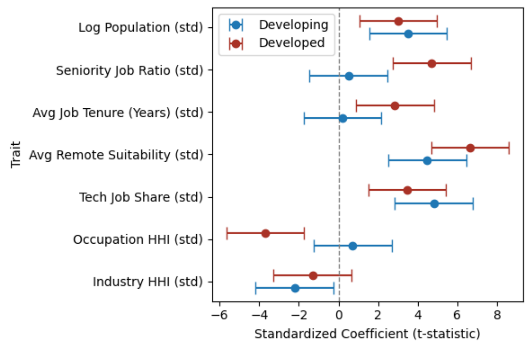

(a) By Development Status

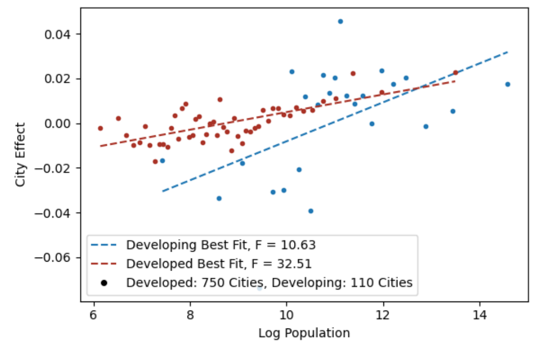

(b) Population

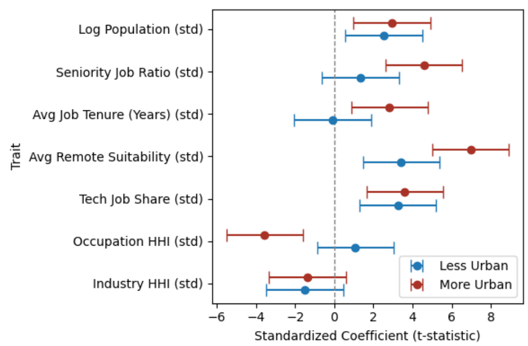

(c) By Urbanization

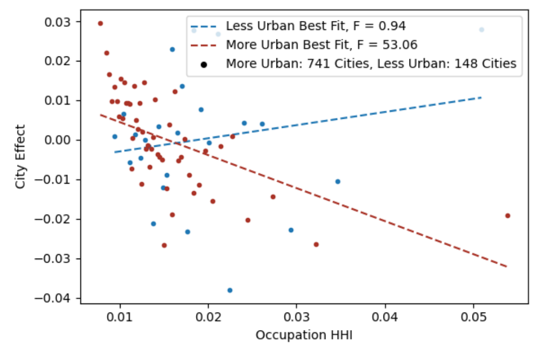

(d) Occupation HHI

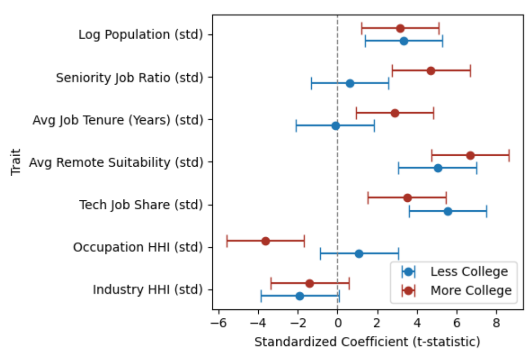

(e) By College Attainment

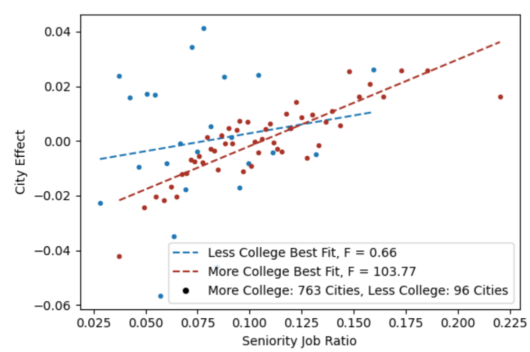

(f) Seniority Ratio

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.

Figure A14: Correlates with City Effects By Country Characteristics

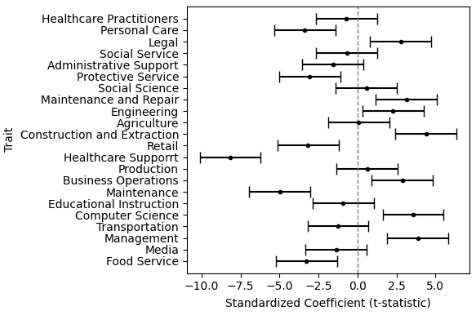

(a) Correlates of City Effects

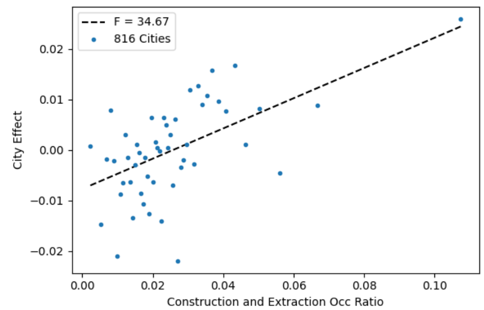

(b) Construction

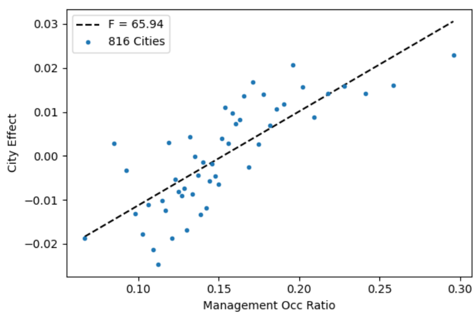

(c) Management

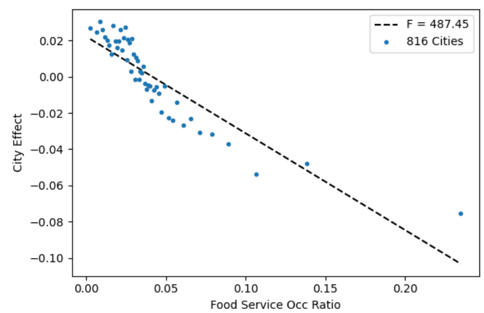

(d) Food Service

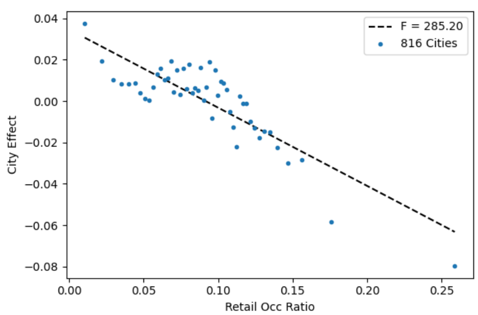

(e) Retail

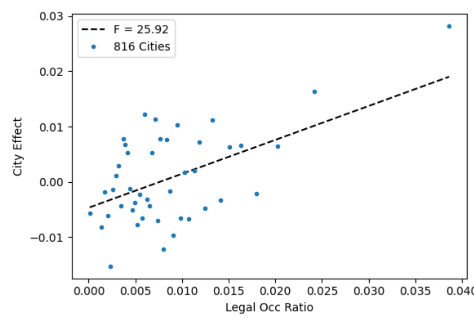

(f) Legal

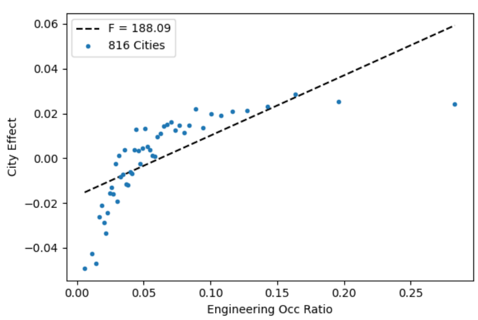

(g) Engineering

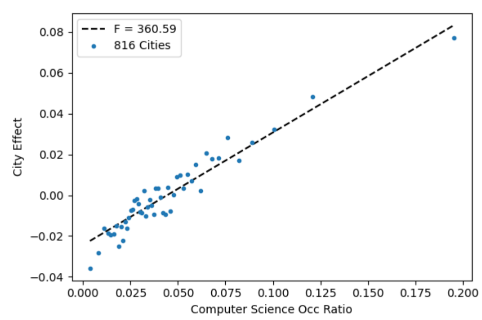

(h) Computer Science

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.

Figure A15: Weather and City Effects

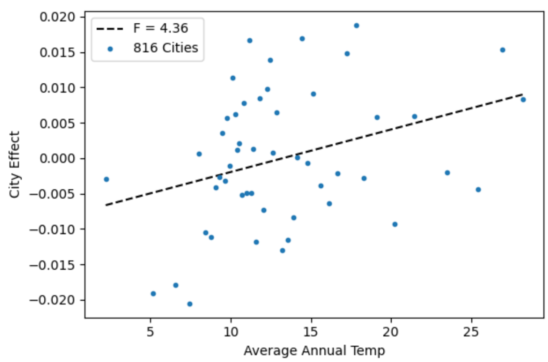

(a) Average Annual Temp

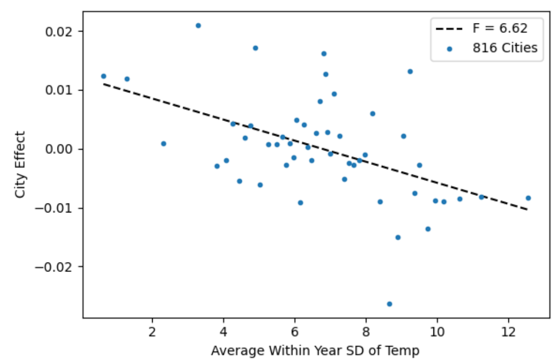

(b) Within Year SD of Temp

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 average annual temperature over the last 20 years. Panel (b) plots the relationship between city effects and the average within-year standard deviation of temperatures over the last 20 years.

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.

Table A3: Metro Grouping Variance Decomposition

	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	1.33	1.92	2.7	1.97	–
Difference due to City	1.22	1.77	2.44	1.88	–
Share due to City	0.92	0.92	0.9	0.95	0.94
Bounded Share					(0.81)

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	0.18	0.29	0.49	0.17	–
Difference due to City	0.08	0.13	0.22	0.07	–
Share due to City	0.43	0.45	0.45	0.41	0.53
Bounded Share					(0.45)

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	0.19	0.32	0.63	0.15	–
Difference due to City	0.13	0.22	0.44	0.01	–
Share due to City	0.68	0.68	0.7	0.05	0.74
Bounded Share					(0.64)

Notes: This table presents the decomposition of wage differences across metro areas 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 bound for the explained share based on Equation 32. To group locations into metro areas, we classified all locations with a population greater than 90,000 as anchor cities, and then grouped non-anchor locations within 30km of anchor cities into the same metro group. We then limited our sample to these metro areas before continuing the analysis.

Table A4: City effects computed using all moves versus just one-time movers (our sample)

	United States (1)	India (2)	All (3)	All (4)
Dependent Variable: City Effect (All Movers)				
City Effect (Our Sample)	0.9782***	0.8287***	0.9671***	0.8851***
	(0.0094)	(0.0194)	(0.0028)	(0.0159)
Country FEs				Yes
R^2	0.9979	0.9964	0.9997	0.9998
Within R^2				0.9935

Notes: Dependent variable was city effect using all movers in the dataset, not just those that move between two cities. Robust standard errors in parenthesis. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

B Bounding the Share of Variance Explained by City Effects

This appendix gives the full imputation-bias model summarized in Section 4.3, derives the lower bound on Ω that we use throughout the paper, and validates it against true wages. The bound uses a reverse-regression argument to construct an upper bound on the bias factor B from any pair of true and imputed firm-effect estimates.

B.1 A model of imputation bias

Recall the true-wage model (9) and the ability decomposition (10). The imputation observes a potentially noisy, possibly position-specific proxy for $X_i\beta$, written $Z_{it} = X_i\beta + \nu_{it}$, where worker ability $X_i\beta$ is time-invariant and ν_{it} is mean-zero measurement noise independent of α_i and $f(i,t)$.[25]

We allow for sorting on both components of ability:

$$\mathbb{E}[X_i\beta \mid f(i,t) = f] = a_X + \phi\,\gamma_f, \tag{17}$$

$$\mathbb{E}[\eta_i \mid f(i,t) = f] = a_\eta + \theta\,\gamma_f, \tag{18}$$

where $\phi \geq 0$ captures sorting on observable ability and $\theta \geq 0$ sorting on unobservable ability. Either can be zero if sorting is weak; both are positive under standard positive assortative matching. To see how these generate bias, consider the firm dummy in the imputation regression. Substituting (9)–(10) into $y_{it} - \lambda Z_{it}$ and taking firm means, using $\mathbb{E}[\varepsilon_{it} \mid f] = \mathbb{E}[\nu_{it} \mid f] = 0$, gives

$$\hat{\gamma}_f^{\text{imp}} = \mathbb{E}[y_{it} - \lambda Z_{it} \mid f] = \gamma_f \;+\; \underbrace{(1-\lambda)\,\mathbb{E}[X_i\beta \mid f]}_{\text{observable ability leakage}} \;+\; \underbrace{\mathbb{E}[\eta_i \mid f]}_{\text{unobservable ability leakage}}.$$

The imputation's firm dummy equals the true firm effect γ_f plus two worker-ability components that the regression fails to assign to the worker side. The first is the share $(1-\lambda)$ of observable ability left behind by attenuation in the coefficient on Z; the second is unobservable ability η, which the imputation cannot observe. Because sorting makes both increasing in γ_f, the firm coefficient loads on γ_f with weight greater than one whenever $\phi > 0$ or $\theta > 0$.

The imputation predicts wages by regressing y on the proxy Z and firm indicators in a training sample. Because Z measures observable ability with error, the OLS coefficient on Z is attenuated toward zero, with attenuation factor λ given by the reliability of the proxy:

$$\hat{\beta}_Z = \lambda, \qquad \lambda = \frac{\mathrm{Var}(X_i\beta)}{\mathrm{Var}(X_i\beta) + \mathrm{Var}(\nu_{it})} \leq 1. \tag{19}$$

The parameter λ captures the share of variation in Z that reflects true ability rather than measurement noise; it equals one when observables are perfectly informative ($\mathrm{Var}(\nu_{it}) = 0$) and falls toward zero as the proxy becomes noisier.[26] Even when $\lambda = 1$ the imputation recovers only $X_i\beta$ and leaves η_i in the residual, so the unobservable-ability channel remains active.

The firm dummies then absorb the wage variation not attributed to Z:

$$\hat{\gamma}_f^{\text{imp}} = \mathbb{E}\big[y_{it} - \lambda Z_{it} \mid f\big] = \gamma_f + (1-\lambda)\,\mathbb{E}[X_i\beta \mid f] + \mathbb{E}[\eta_i \mid f]. \tag{20}$$

[25] Letting the proxy noise ν_{it} vary across a worker's positions, while $X_i\beta$ is fixed, is what makes the differenced noise term in (23) mean zero.

[26] With firm indicators included in the imputation regression, the operative object is the *within-firm* reliability ratio $\mathrm{Var}_{\text{within}}(X_i\beta)/\big[\mathrm{Var}_{\text{within}}(X_i\beta) + \mathrm{Var}(\nu_{it})\big]$; under sorting this is weakly smaller than the unconditional ratio in (19), which we report for transparency. The distinction does not affect the empirical bound below, which is computed from the reverse-regression slope β_R directly rather than from λ.

Substituting (17) and (18) and collecting constants,

$$\hat{\gamma}_f^{\text{imp}} = \big(1 + (1-\lambda)\phi + \theta\big)\,\gamma_f + \text{const} = (1+B)\,\gamma_f + \text{const}, \tag{21}$$

with $B \equiv (1-\lambda)\phi + \theta$ as in Equation (11). The imputed wage therefore takes the form

$$\hat{y}_{it} = \lambda Z_{it} + (1+B)\,\gamma_{f(i,t)} + \xi_{it}, \tag{22}$$

where ξ_{it} is a mean-zero imputation residual orthogonal to (X, f).

Estimating Equation (4) on $\hat{y}_{it}$ identifies γ from within-worker wage changes around moves. For a worker who switches from firm A to firm B,

$$\hat{y}_{it} - \hat{y}_{i,t-1} = \lambda(\nu_{it} - \nu_{i,t-1}) + (1+B)\,(\gamma_B - \gamma_A). \tag{23}$$

The first term is mean-zero noise; the systematic component is the second. AKM on $\hat{y}$ therefore identifies firm effects scaled by $(1+B)$, which is Equation (12) in the main text.

Equation (11) shows that imputation generates upward bias through two channels, both rooted in worker–firm sorting. The first channel, $(1-\lambda)\phi$, combines imperfect recovery of observable ability $(1-\lambda)$ with sorting on observables (ϕ): when the imputation under-uses available covariates, firm dummies absorb leftover worker-side variation correlated with firm productivity through sorting. Richer imputation (higher λ) shrinks this channel; no sorting on observables ($\phi = 0$) eliminates it. The second channel, θ, captures sorting on unobserved ability: the imputer attributes co-worker unobservable ability to the firm component of its prediction, and better imputation cannot reduce this without observing η directly. Neither channel reflects bias in AKM itself; AKM on true wages remains unbiased under arbitrary positive assortative matching. The bias arises because the imputation cannot recover all of α_i, and firm dummies absorb the missing portion in proportion to how strongly it sorts across firms.

A second feature of the framework is that AKM on $\hat{y}$ is insensitive to additional imputation noise uncorrelated with (X, f). Identification of $\tilde{\gamma}$ from movers' wage changes differences out time-invariant individual factors, so mean-zero imputation noise widens confidence intervals but does not introduce systematic bias. Imputation quality enters Equation (11) only through λ, the systematic attenuation in how the imputation uses observables, not through the variance of imputation residuals.

B.2 From firm effects to the variance share

Aggregating the firm-level relation (21) to the city level using Equation (5) gives the city-effect scaling (13), $\tilde{\Gamma}_j = (1+B)\,\Gamma_j + \text{const}$. The variance share $\Omega = \text{Cov}(\bar{y}_j, \Gamma_j)/\text{Var}(\bar{y}_j)$ inherits the same factor under the following restriction on how worker ability is distributed across cities.

Assumption. City-mean worker ability is an affine function of the city effect,

$$\mathbb{E}[\alpha_i \mid \text{city } j] = a + (\phi + \theta)\,\Gamma_j, \tag{24}$$

equivalently, the firm-level sorting relations (17)–(18) aggregate to the city level with no component orthogonal to Γ_j.[27]

[27] Time-varying controls $\bar{x}_j\beta$, being functions of observables, carry through analogously and are absorbed into the constant below.

Under this assumption, true and imputed city-mean wages are both affine in Γ_j with the same slope. For true wages, $\bar{y}_j = \text{const} + (1+\phi+\theta)\,\Gamma_j$. For imputed wages, averaging (22) within a city and using $\mathbb{E}[Z_{it} \mid j] = \mathbb{E}[X_i\beta \mid j] = a_X + \phi\,\Gamma_j$ gives

$$\bar{\hat{y}}_j = \lambda\,\mathbb{E}[X_i\beta \mid j] + (1+B)\,\Gamma_j + \text{const} = \text{const} + \big(\lambda\phi + 1 + B\big)\Gamma_j = \text{const} + (1+\phi+\theta)\,\Gamma_j,$$

where the last equality uses $B = (1-\lambda)\phi+\theta$. Both city-mean wages thus share the slope $(1+\phi+\theta)$, while estimated city effects scale by $(1+B)$. Hence

$$\hat{\Omega} = \frac{\operatorname{Cov}\big(\bar{\hat{y}}_j,\,\tilde{\Gamma}_j\big)}{\operatorname{Var}\big(\bar{\hat{y}}_j\big)} = (1+B)\,\frac{\operatorname{Cov}(\bar{y}_j,\Gamma_j)}{\operatorname{Var}(\bar{y}_j)} = (1+B)\,\Omega,$$

which is Equation (14). The assumption rules out city-level differences in worker composition unrelated to the city effect; when it holds, the inflation in estimated city effects carries over to the variance share.

B.3 Setup

The city-effect scaling $\tilde{\Gamma}_j = (1+B)\,\Gamma_j + \text{const}$ in Equation (13) holds for population objects. In practice both $\tilde{\Gamma}_j$ and Γ_j are estimated. Writing the AKM estimates as the true population objects plus sampling error:

$$\hat{\Gamma}_j^{\text{true}} = \Gamma_j + \mu_j, \tag{25}$$

$$\hat{\tilde{\Gamma}}_j = (1+B)\,\Gamma_j + \tilde{\nu}_j, \tag{26}$$

where μ_j collects AKM sampling error on the true-wage regression and $\tilde{\nu}_j$ collects sampling error on the imputed-wage regression as well as any deviation from the exact constant-B relationship. We impose the standard regularity conditions:

$$\mathbb{E}[\mu_j] = \mathbb{E}[\tilde{\nu}_j] = 0, \quad \operatorname{Cov}(\Gamma_j,\mu_j) = \operatorname{Cov}(\Gamma_j,\tilde{\nu}_j) = 0, \quad \operatorname{Cov}(\mu_j,\tilde{\nu}_j) = 0. \tag{27}$$

The first three conditions are standard in the AKM literature. The fourth, uncorrelated noise across the two AKM regressions, is a stronger restriction that we revisit below.

B.4 Frisch reverse-regression bounds

The forward OLS slope from regressing $\hat{\tilde{\Gamma}}_j$ on $\hat{\Gamma}_j^{\text{true}}$ is, under (25)–(27),

$$\beta_F \;\equiv\; \frac{\operatorname{Cov}\big(\hat{\tilde{\Gamma}}_j,\,\hat{\Gamma}_j^{\text{true}}\big)}{\operatorname{Var}\big(\hat{\Gamma}_j^{\text{true}}\big)} \;=\; (1+B)\,\frac{\operatorname{Var}(\Gamma_j)}{\operatorname{Var}(\Gamma_j)+\operatorname{Var}(\mu_j)}. \tag{28}$$

Since $\operatorname{Var}(\mu_j) \geq 0$, the ratio on the right is at most one, so $\beta_F \leq (1+B)$. The forward slope is the standard attenuation of $(1+B)$ by the reliability of $\hat{\Gamma}_j^{\text{true}}$ as a measure of Γ_j.

The reverse OLS slope from regressing $\hat{\Gamma}_j^{\text{true}}$ on $\hat{\tilde{\Gamma}}_j$ is

$$\beta_R \;\equiv\; \frac{\operatorname{Cov}\big(\hat{\Gamma}_j^{\text{true}},\,\hat{\tilde{\Gamma}}_j\big)}{\operatorname{Var}\big(\hat{\tilde{\Gamma}}_j\big)} \;=\; \frac{(1+B)\operatorname{Var}(\Gamma_j)}{(1+B)^2\operatorname{Var}(\Gamma_j)+\operatorname{Var}(\tilde{\nu}_j)}. \tag{29}$$

Inverting,

$$\frac{1}{\beta_R} \;=\; (1+B) + \frac{\mathrm{Var}(\tilde{\nu}_j)}{(1+B)\,\mathrm{Var}(\Gamma_j)} \;\geq\; (1+B). \tag{30}$$

Combining (28) and (30),

$$\beta_F \;\leq\; (1+B) \;\leq\; \frac{1}{\beta_R}. \tag{31}$$

This is the interval for $(1+B)$, requiring only the standard AKM noise structure and the additional assumption that sampling errors across the two regressions are uncorrelated.

B.5 Implied lower bound on Ω

By Equation (14), $\Omega = \hat{\Omega}/(1+B)$. Any upper bound on $(1+B)$ delivers a lower bound on Ω. Using the upper end of the Frisch interval,

$$\Omega \;\geq\; \frac{\hat{\Omega}}{1/\beta_R} \;=\; \beta_R \cdot \hat{\Omega}. \tag{32}$$

The bound is a function of two quantities: the reverse OLS slope β_R from comparing imputed and true firm effects, and the share of variance $\hat{\Omega}$ from the imputed-wage AKM in the country of interest.

B.6 Empirical estimates and validation

We quantify the bias factor B and the implied bound on Ω using the matched employer–employee data from Italy introduced in Section 4.3, where we observe true wages and can construct imputed wages following the structure of the Revelio procedure. We construct two wage imputations that differ in the richness of observable information used. The more detailed imputation predicts wages using average wages among full-time workers in the same firm, tenure class, qualification, and year; based on Revelio's documentation, this uses information similar to the global imputation. The less detailed imputation uses only firm and qualification. Both perform well at the individual level: regressions of imputed wages on true wages yield $R^2 = 0.876$ and $R^2 = 0.766$, respectively. Figure A16a plots the more detailed imputed wages against true wages, showing a tight relationship between the two measures.

For each imputation we estimate firm effects by AKM on imputed wages, aggregate them to city effects using Equation (5), and compute the share $\hat{\Omega}$. Using true wages, $\Omega = 0.584$. The more detailed imputation gives $\hat{\Omega}_{\text{more}} = 0.658$, implying $B_{\text{more}} = \hat{\Omega}_{\text{more}}/\Omega - 1 = 0.127$; the less detailed imputation gives $\hat{\Omega}_{\text{less}} = 0.758$ and $B_{\text{less}} = 0.299$. As predicted by the model, the more detailed imputation produces a smaller bias factor because it raises λ and shrinks the imputation-quality channel $(1-\lambda)\phi$. The 12.7% overstatement under the more detailed imputation corresponds to 7.4 percentage points of the share attributable to city effects. We find no evidence that the excess variance is correlated with location-level observables such as the number of firms, the number of workers, or average true wages.

Applying the bound in Equation (32), we obtain a lower bound of $\Omega = 0.568$ for the more detailed imputation and 0.529 for the less detailed imputation. The bound is conservative: it understates the true Ω by 2.7% and by 9.3%, respectively.

A second concern is that Revelio's global imputation may be less accurate than our Italian construction, in which case the Italian-calibrated bound parameter β_R may not extend to other countries. To assess this, we progressively add mean-zero Gaussian noise to the more detailed imputed wages and recompute $\hat{\Omega}$, defining the noise standard deviation as a percentage of the standard deviation of imputed wages. Figure A16c plots added noise against $\hat{\Omega}/\Omega - 1$, the implied B. The relationship is weak below one standard deviation of imputed wages: adding noise of one standard deviation raises the estimated bias by only about 3 percentage points. This insensitivity to plausible degradation in imputation quality supports the use of the Italian bound as a conservative global adjustment.

For the global application we use the Italian β_R from the more detailed imputation as our bound parameter, under the assumption that Revelio's global imputation does not produce a larger B than our Italian construction. Two arguments support this. First, Revelio's imputation uses substantially more covariates than the firm × tenure × qualification × year averaging we implement in Italy—job titles grouped into 1,500 categories, granular tenure and seniority measures, geographic and economic indicators, and company identifiers—and a richer covariate set implies a higher λ and therefore a smaller imputation-quality channel. Second, the noise stress test shows that degrading the Italian imputation changes the estimated bias only modestly. We report this bound alongside our headline estimates of Ω throughout the paper.

Figure A16: Bias in city effect estimation arising from wage imputation

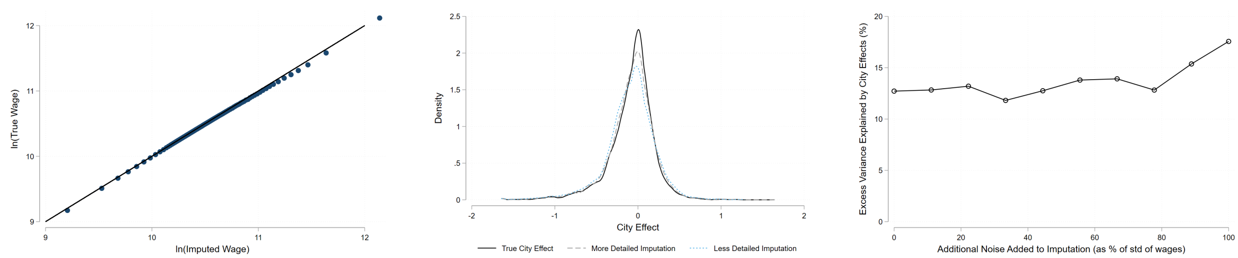

(a) Binned scatter of true vs. imputed wages

(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.

B.7 When the bound holds, when it fails

The interval in (31) is exact under Equation (27). Two restrictions on the assumptions are worth flagging.

The condition $\text{Cov}(\mu_j, \tilde{\nu}_j) = 0$ is restrictive in practice. Both noise terms arise from AKM regressions on the same firm network and overlapping mover patterns, so they are likely positively correlated. If $\text{Cov}(\mu_j, \tilde{\nu}_j) > 0$, the forward slope becomes

$$\beta_F \;=\; (1+B)\,\frac{\text{Var}(\Gamma_j)}{\text{Var}(\Gamma_j)+\text{Var}(\mu_j)} \;+\; \frac{\text{Cov}(\mu_j, \tilde{\nu}_j)}{\text{Var}(\Gamma_j)+\text{Var}(\mu_j)}, \tag{33}$$

and may exceed $(1+B)$. The lower endpoint of the Frisch interval can be violated. The upper endpoint remains valid as long as

$$\text{Cov}(\mu_j, \tilde{\nu}_j) \;\leq\; \frac{\text{Var}(\tilde{\nu}_j)}{1+B}, \tag{34}$$

which holds unless the two noises are nearly perfectly correlated. Since our application uses only the upper endpoint $1/\beta_R$ for the bound on Ω, we take this as the relevant condition.

The bound also rests on the linearity of $\tilde{\Gamma}_j$ in Γ_j. If B varies systematically across cities, for example if sorting on observables and unobservables is stronger in larger or more urbanized cities, then no single multiplicative factor describes the relationship, and the interval bounds some weighted average rather than a uniform bias.