Location, Location, Location

David Card, UC Berkeley

Jesse Rothstein, UC Berkeley

Moises Yi, US Census Bureau

August 2023

ABSTRACT

We use data from the Longitudinal Employer-Household Dynamics program to study the causal effects of location on earnings. Starting from a model with employer and employee fixed effects, we estimate the average earnings premiums associated with jobs in different commuting zones (CZs) and different CZ-industry pairs. About half of the variation in mean wages across CZs is attributable to differences in worker ability (as measured by their fixed effects); the other half is attributable to place effects. We show that the place effects from a richly specified cross sectional wage model overstate the causal effects of place (due to unobserved worker ability), while those from a model that simply adds person fixed effects understate the causal effects (due to unobserved heterogeneity in the premiums paid by different firms in the same CZ). Local industry agglomerations are associated with higher wages, but overall differences in industry composition and in CZ-specific returns to industries explain only a small fraction of average place effects. Estimating separate place effects for college and non-college workers, we find that the college wage gap is bigger in larger and higher-wage places, but that two-thirds of this variation is attributable to differences in the relative skills of the two groups in different places. Most of the remaining variation reflects the enhanced sorting of more educated workers to higher-paying industries in larger and higher-wage CZs. Finally, we find that local housing costs at least fully offset local pay premiums, implying that workers who move to larger CZs have no higher net-of-housing consumption.

*We thank David Albouy, Wolfgang Dauth, Donald Davis, Rebecca Diamond, Enrico Moretti, Matthew Kahn, Pat Kline, Jorge De la Roca, Alice Wu and seminar participants at UC Berkeley, Columbia, University of Florida, USC, Carnegie Mellon University, George Washington University, the Federal Reserve Board of Governors, the Urban Institute, Centre de Recerca en Economia Internacional, SOLE, CESifo Labor, NY Fed, Philadelphia Fed, Mid-Atlantic Trade Workshop, and NBER Labor Studies for many helpful comments and suggestions. We are also grateful to Nicole Gandre and Richard Jin for outstanding research assistance. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 6000266: CBDRB-FY21-122, CBDRB-FY21-149, CBDRB-FY21-288, CBDRB-FY22-021, CBDRB-FY22-065, CBDRB-FY22-135, CBDRB-FY23-0293, CBDRB-FY23-0395, CBDRB-FY23-0410).

There are large geographic differences in earnings in the U.S. and other countries. Bigger cities tend to have higher wages (Behrens, et al., 2014; Eeckhout et al., 2014; Butts et al. 2023) and higher returns to education (Autor, 2019; Davis and Dingel, 2019), but there are also wide disparities among cities of similar size. The source of these differences is a perennial issue in economic geography.

In a simple Roback (1982)-style model of spatial equilibrium, geographic differences in nominal wages mirror productivity differences; otherwise employers will move.[1] Three broad explanations have been offered for local productivity differences. The first is systematic sorting of higher-skilled people to preferred places (e.g., Behrens et al., 2014). A second is the presence of productive industries (e.g., Ellison and Glaeser, 1997; Rosenthal and Strange, 2004). The third is causal effects of the places themselves, arising from endogenous externalities of population density or human capital (Ciccone and Hall, 1996; Duranton and Puga, 2004; Glaeser and Gottlieb, 2009; Diamond, 2016), or from exogenous factors like geography or climate.

These explanations have sharply different predictions for the impacts of worker mobility. A sorting-based explanation implies that mobility has no effect on earnings. In contrast, theories based on industry composition or place-based factors imply that workers become more productive, and earn higher wages, when they move to larger or higher-wage places. Different explanations also have distinct implications for policy. Industry-based explanations imply that regions can benefit from clusters of high-wage industries, providing a rationale for tax subsidies (e.g., Greenstone, Hornbeck, and Moretti 2010). Likewise, endogenous externalities suggest a role for policies to increase local population or attract highly skilled workers (Moretti 2004a,b). Theories based on local factors like climate leave less room for policy, but point to the importance of increasing housing supply in highly productive places.

[1] In this model, a similar no arbitrage condition for households relates real, local cost adjusted wages to local consumption amenities. We discuss regional price differences and real wages in the final section of this paper.

Despite several decades of research there is little consensus on the relative importance of these three explanations, or even on the simple question of whether movers experience systematic wage changes that are correlated with conventionally estimated place effects. A seminal study by Glaeser and Maré (2001) found a mixed pattern of evidence on the effects of moving into or out of metropolitan areas, depending on the data set and direction of the move.[2] Subsequent studies using large administrative data sets from outside the U.S. (Combes et al., 2008; Gibbons et al., 2014; de la Roca and Puga, 2017; Dauth et al., 2022) have found evidence of skill-based sorting *and* place effects that both contribute to observed earnings differentials.

In this paper we use data from the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program to measure the average pay differences associated with working in different commuting zones (CZs).[3] Our approach begins at the level of individual establishments. Building on the seminal paper of Abowd, Kramarz, and Margolis (1999; hereafter *AKM*) we fit models for quarterly log earnings that include person fixed effects, controls for age and time, and establishment fixed effects. We then average the estimated effects across establishments in each place to obtain CZ wage premiums. These represent the gains or losses in pay associated with a move from the average establishment in one CZ to the average establishment in another.

In a first methodological contribution we show that two-way fixed effects models that include person and place effects (as in Glaeser and Maré [2001], Combes et al. [2008], and earlier versions of this paper [Card et al. 2022]) yield biased estimates of place effects.[4] The problem is that moves

[2] Appendix Figure 1 presents a visual summary of Glaeser and Maré's main results. They used the Panel Study of Income Dynamics and the National Longitudinal Study of Youth, both of which have relatively small samples. A later study by Gould (2007) concluded that gains for migration are confined to high-education workers.
[3] CZs are meant to capture regions within which workers commute. Unlike metropolitan statistical areas (MSAs), they cover the entire country, including both urban and rural areas. The locations of some workplaces in the LEHD are imputed. As discussed below, however, we find that the effects of any imputation errors are negligible.
[4] In a companion paper focusing on pay differences across industries (Card, Rothstein and Yi, 2023) we show that the same misspecification issues arise in traditional industry mover designs, such as the one proposed by Krueger and Summers (1988).

between geographic areas, which identify the place effects in such models, are associated with systematic changes in the relative pay premiums of the establishments within the origin and destination CZs. Workers who move from a lower-paying CZ to a higher-paying CZ tend to come from workplaces that offered above-average wage premiums in their origin city, and move to workplaces with below-average premiums in their destination city. A symmetric shift occurs for workers who move in the opposite direction. This pattern of selective mobility between establishments at different levels of the local job ladder -- which we call a "hierarchy effect" -- holds true even accounting for changes in industry of CZ movers. Specifications that fail to account for this pattern yield attenuated estimates of the pay premiums for different CZs. Our AKM-based methodology avoids this attenuation.

We use event-study style comparisons of earnings before and after a move to show that between-CZ mobility is approximately exogenous with respect to the residuals in our underlying AKM model. This finding parallels recent evidence for firm-to-firm mobility in other countries, and suggests that the place effects from our approach are approximately unbiased. We also present, in the appendix, extensive validation evidence regarding the AKM model, providing what we believe is the first systematic assessment of the fit of the model to a U.S.-wide sample.

A limitation of the conventional AKM framework is the assumption that only the current employer (and current city) matter for pay. In an influential study using administrative data from Spain, De la Roca and Puga (2017) find that young workers who spend time in larger cities earn more later, even when they move to smaller cities. We estimate augmented AKM specifications that incorporate dynamic effects of spending time in large, high-wage CZs. The size of the U.S. labor market is an advantage here – the ten largest U.S. CZs have a similar combined population share as Madrid and Barcelona, the two largest cities in Spain. We find statistically significant but modest-sized dynamic effects that are consistent with the mechanism postulated by De la Roca and Puga (2017). Accounting for these effects, however, has very little impact on our estimates of average CZ pay differences. Accordingly, most of our analysis emphasizes a simpler model without dynamics.

Our second contribution is to identify the effects of skill-based sorting on observed earnings differences across CZs. Only 50% of the variation in mean earnings across CZs is attributable to place effects. The other half is due to worker sorting. As in France (Combes et al., 2008), Spain (de la Roca and Puga, 2017), and Germany (Dauth et al., 2022), we find that higher earnings capacity workers are more likely to live in high-wage CZs, magnifying the inequality in average earnings differences across CZs. Similar conclusions hold with respect to CZ size: Half of the observed earnings premium associated with working in a larger CZ reflects worker sorting. We also find that larger places have more dispersion in skills (Eeckhout et al., 2014), and stronger assortative matching between high skilled workers and high-premium firms (Dauth et al., 2022).

Our third contribution is to characterize the role of industry specialization in CZ wage premia. We ask how important are differences in the shares of high-wage industries, and in CZ-specific returns to specific industries, to the variation in estimated place effects. We construct place-by-industry effects by averaging the firm effects for all firms in a CZ-industry cell. We then develop a simple decomposition of CZ-by-industry pay premiums into a combination of place effects, industry effects, and match (or interaction) effects that reflect CZ-specific returns to different industries. We find significant "main effects" of places and industries, but only small match effects – i.e., close to additive separability in the contributions of place and industry.[5]

Although there is some variation in the share of high wage industries in different cities, we find that the overall contribution of industry composition to average pay differences is small. Consistent with the presence of agglomeration effects, we also find that the excess pay premium for a given industry in

[5] This approximate additive separability allows us to separate the study of place and industry effects. This paper studies place effects. A companion paper (Card et al. 2023) studies industry effects.

a CZ (i.e., the match effect) is higher when its local employment share is higher, but again the impact of these interaction effects on average CZ pay premiums is quite modest.

Our fourth contribution is to explore locational differences in the pay gap between different education groups. In cross-sectional models, college-educated workers receive larger returns from living in larger or higher-wage cities, a fact that has received much attention in the recent literature (e.g., Autor, 2019; Davis and Dingel, 2019, 2020). But such comparisons are confounded by differences in unobserved skills within education groups. To address this concern we estimate separate CZ premiums for college and non-college workers. We find that the skills of the college-educated workforce are higher in larger and higher-wage places (consistent with Diamond, 2016). This differential sorting explains approximately two-thirds of the observed correlation between CZ size or CZ average earnings and the place-specific gap between college and high school wages.[6] Much of the remaining third is attributable to enhanced sorting of college-educated workers to high-wage industries in larger and higher-wage CZs. Only about a tenth reflects differences in the local wage premia for high- versus low-education workers.

We conclude by examining housing cost differences across CZs, and the implied differences in real (cost-of-living adjusted) earnings. We find large elasticities of housing costs with respect to CZ mean wages and log size – large enough to at least offset the corresponding effects on nominal earnings. Thus, movements to larger or higher-wage locations yield reductions in real income. In a Roback-style model this pattern suggests that more productive places have higher consumption amenities that offset their higher cost of living.

Our work is related to three main literatures. The first is a set of recent studies that use large administrative data sets from outside the U.S. and fixed effects specifications to separate the effects of

[6] Our data cover 2010-2018, and do not allow us to examine how either sorting or CZ effects have changed over recent decades, as suggested by Autor (2019) and Butts et al. (2023).

place from the non-random sorting of workers.[7] Combes et al. (2008) estimate models on French data that include fixed effects for workers, employment areas, and industry. De la Roca and Puga (2017) estimate models for Spain that include fixed effects for workers, urban areas (UA's), and measures of cumulative work experience in larger UAs. Both studies conclude that about half of the locational premium for larger areas is attributable to sorting of higher-earning workers. Dauth et al. (2022) estimate models for Germany that include fixed effects for workers and *establishments* and conclude that up to three quarters of the wage premium for working in a larger city reflects worker quality and enhanced sorting between workers and firms.

We contribute to this literature by providing the first estimates from administrative data for the U.S., and by carefully evaluating the underlying assumptions in our own and previous specifications. We show that the specifications with person and place effects used by most previous studies (with the notable exception of Dauth et al. [2022]) are likely to yield attenuated estimates of causal place effects due to differences in the pay premiums at different employers in the same local market that are correlated with the direction of between-place moves.

Second, we relate to the large literature in urban economics on market size elasticities and the returns to agglomeration (Rosenthal and Strange, 2004; Baum-Snow and Pavan, 2012; Behrens et al., 2014; Eekhout et al., 2014; Butts et al. 2023). A related literature considers the impact of high-wage employers or industries on local economic development (Greenstone, Hornbeck, and Moretti 2010; Kline and Moretti 2014) or worker location choices (Diamond, 2016). We contribute to this literature by showing that there is substantial sorting of higher-skilled workers to larger, higher-wage CZs. Controlling for worker skills and national industry wage premiums, we find that local industry structure, industry-based agglomerations, and match effects explain only small additional shares of CZ wage differences.

[7] A related set of papers use similar strategies to study intergenerational mobility, health costs, and other topics. See, e.g., Chetty and Hendren (2018a,b) and Finkelstein, Gentzkow, and Williams (2016, 2021).

However, the sorting of higher-skilled workers to higher-paying firms and industries is enhanced in larger CZs, suggesting an important benefit to local agglomeration (Dauth et al., 2022).

Finally, our work is related to the growing literature that examines firm-specific pay differences. Studies in this vein focus on *firms* or *plants* as the unit of analysis and often ignore the role of geography. We show that location is an important component of the pay premiums offered by different establishments in the U.S. that is approximately orthogonal to the industry component described in our companion paper (Card et al., 2023).

II. Geographic earnings premiums in the U.S.

As a starting point for our analysis it is useful to review four key facts about the geographic variation in pay in the U.S. These are generally well known in the urban economics literature; we follow previous authors by using data from the American Community Survey (ACS) – in our case pooling information from the 2010-2018 surveys to parallel the LEHD samples we use below – and focusing on commuting zones (CZs) as our basic unit of geography.[8]

First, mean wages vary widely across CZs. Across nearly 700 CZs that can be identified in the ACS, the (population-weighted) standard deviation of CZ-specific mean log hourly wages is 0.14, while the interquartile range is 0.20. Wages in San Francisco, San Jose, and Washington DC are about 30 log points above the national mean, while wages in Brownsville and Laredo are about 30 log points below.

Second, only a modest share of this variation is explained by differences in observed characteristics of the workers. We used a model with about 150 individual covariates (including flexible functions of experience, education, gender, and race/ethnicity; controls for region of origin and time in the country for immigrants, and controls for college major) as well as detailed industry controls (>260

[8] Commuting zones are intended to approximate integrated labor market areas, with each CZ comprised of one or more complete counties. We use the 1990 vintage of CZ definitions.

categories) and CZ dummies to derive composition-adjusted CZ effects. The standard deviation of the resulting CZ effects is 0.106 – about 75% as large as the variation in unadjusted mean wages. In Appendix Figure 2 we plot the adjusted CZ effects against the mean wages in each CZ. The slope is 0.692, suggesting that only about 30% of the variance of CZ wages can be explained by even a rich set of observed worker and job characteristics.[9]

Third, mean wages are higher in larger CZs, with an elasticity of about 0.068 (standard error = 0.010). Again, only a small share of this elasticity is explained by observed worker characteristics. The elasticity of composition-adjusted CZ effects (from the model described above) with respect to CZ size is 0.056 (std. error = 0.005) – about 80% as large as the elasticity of unadjusted wages. These relationships are illustrated in Appendix Figure 3.[10]

A fourth fact – highlighted in recent work by Autor (2019) and Davis and Dingel (2019; 2020) – is that the pay of better-educated workers rises more in larger or higher mean-wage CZs. For example, the CZ size elasticity of mean log wages for workers with exactly 16 years of schooling is 0.071 (std. error = 0.007), compared to only 0.030 (std. error = 0.004) for workers with 12 years of schooling. If we fit separate versions of the model described above for the two groups, and regress the adjusted CZ wage effects for the two subgroups on the estimated CZ effects for all workers, we obtain coefficients of 0.906 for workers with 12 years of schooling and 1.209 for those with 16 years of schooling. The subgroup effects are nearly perfectly correlated with each other and with the pooled CZ effects, indicating that the stakes are higher for more educated workers, but the ranking of CZs is the same across education.

[9] We can use a similar model (omitting the industry controls) to examine geographic variation in employment rates or annual hours. While CZ mean wages are positively correlated with CZ mean labor supply, this correlation falls to zero when worker characteristics are controlled. See Appendix Table 2.

[10] In Appendix B and Appendix Table 1 we explore differences in the size elasticities of hourly wages and annual wages using ACS data. The size elasticity of unadjusted annual earnings is a little larger than the size elasticity of hourly wages, but when we adjust for worker characteristics these gaps disappear.

A critical limitation of cross-sectional data sets like the ACS is that there is no way to control for unmeasured worker skills. Thus, there is no way to know how much of the variation in mean wages across CZs can be explained by skills that are observed by employers and rewarded in the market, but unmeasured in conventional data sets. Likewise, it is unclear how much of the size elasticity of mean wages is driven by the sorting of workers with higher levels of unobserved skills to larger cities, or how much of the higher return to education in larger cities is the result of unobserved skill gaps between college and high school workers that vary with city size.

In the remainder of the paper we use LEHD panel data that follow workers as they move across firms and CZs. These data allow us both to separate permanent earnings differences among workers from any causal place effects, and to address potential selectivity of the firms at which movers are employed within their origin and destination CZs. We start with a brief overview of the LEHD data before describing our approach in more detail.

III. Longitudinal Earnings Data from LEHD

LEHD data are derived from quarterly earnings reports provided by employers to state unemployment insurance agencies, which are then assembled by the Census Bureau into a national data set. The core data include total wages paid by a given employer to each worker in a quarter and a few characteristics of workers and establishments, including industry and location (discussed below). These are supplemented with information on workers and employers collected from other sources (e.g., decennial census and ACS files, linked at the individual level; see Abowd et al., 2009). The LEHD covers about 95% of private sector employment, as well as state and local government employees, but excludes federal employees, members of the armed services, and self-employed workers. From 2010 forward it includes data from all 50 states.

Unfortunately, the LEHD has no information on job start/end dates or hours of work. To help screen out part-time jobs and newly ending or starting jobs we exclude person-employer-quarters (PEQs) where an individual had multiple jobs, or earned less than $3,800 (roughly the earnings from a full-time job at the federal minimum wage). We also exclude all *transitional* quarters (the first and last quarter of any person-employer spell). We then select PEQs for workers age 22-62 from 2010Q1 to 2018Q2. We drop workers who are observed for less than 8 quarters in this period, and PEQs with unknown industry and/or establishment location. We further restrict to workers in the largest connected set of the worker-establishment graph. This connected set includes more than 98% of all PEQs that otherwise meet our sample criteria.

We use information in the LEHD to assign establishments to industries and CZs.[11] Our primary analyses include all 741 CZs in the country. Our investigation of industry effects considers both two-digit and four-digit NAICS industry classifications (with 20 and 312 categories, respectively). The size-weighted correlation between CZ-mean log quarterly earnings from our LEHD sample and CZ-mean log hourly wages from our ACS sample is 0.94.

We examine two subsamples of our main LEHD sample. In some analyses we limit attention to the approximately 15% of workers who can be linked to education information from the ACS (2001-2017). We divide them into more educated (some college or more) and a less educated (high school or less) subgroups. In other analyses we focus on wage dynamics around mobility events (i.e., changes in CZ and/or firm). Here, we limit attention to workers who move between CZs or firms only once in our sample, with stable jobs at the same firm for at least five consecutive quarters before and after the

[11] A potential concern is that the LEHD does not measure directly at which establishment a worker is employed when the firm has multiple establishments in the state. We rely on the LEHD imputation of this establishment, based on the distance between the worker's home and establishment locations. In many cases, all plausible establishments are in the same CZ and industry; we show in Appendix Table 4 that our basic decomposition is unchanged when we limit to these cases.

switch. Because many moves involve periods of non-employment, we allow up to four quarters of non-employment between the origin CZ (or firm) and the destination.

Table 1 presents some characteristics of our LEHD samples, including the fraction observed in different numbers of CZs and industries over the sample period. The first column presents results for our full estimation sample. Columns 2-3 classify this sample by whether people are observed in multiple CZs. CZ movers are younger and more likely to be male than stayers, but have similar average earnings. Finally, columns 4 and 5 summarize our event study samples of people with exactly one move between firms (column 4) or CZs (column 5). As explained above, these samples are selected to have stable employment before and after their moves, so they have somewhat higher average earnings than the full sample.

An initial look at the impacts of mobility

We use our event study sample to conduct an initial descriptive analysis of earnings changes associated with moves between CZs. We construct an adjusted earnings measure as the residual from a regression of quarterly earnings on time effects and a polynomial in age. Then, following Card, Heining, and Kline (2013) (hereafter CHK), we classify CZs into quartiles based on average earnings, yielding 16 origin/destination pairs. Figure 1 plots the means of adjusted wages by quarter relative to the move for the subsets of movers originating from CZs in the top and bottom quartiles.[12]

The figure shows that pre-move earnings are quite stable, with no sign of trends or shocks that predict mobility. Earnings then change immediately following moves, with patterns for different origin/destination groups that suggest a causal effect of places. People who move from bottom-quartile to top-quartile CZs tend to see earnings increases, while those who move from top- to bottom-quartile CZs tend to see declines. The identity of the destination CZ also helps predict the level of pre-move

[12] Recall that we exclude transitional quarters and allow for up to four quarters of non-employment between them. Thus, there may be as many as six quarters between the last observation in the origin CZ (labeled -1) and the first one in the destination (labeled +1).

earnings for workers from the same origin group. People from origin quartile 1 who will move to a quartile 4 CZ, for example, earn more before the move than those who will move to other quartiles. This pattern is consistent with dynamic sorting based on the permanent component of a worker's ability: i.e., within a given origin CZ, higher skilled workers are more likely to move up the locational ladder than their lower-skilled neighbors. A final notable feature of Figure 1 is that earnings tend to grow slightly in first few quarters after arrival in a new CZ.[13]

IV. Methods

In this section we present our methodology for studying locational wage differences. We begin by presenting our person-and-firm fixed effects model and discussing some of the key specification issues in this model, including the potential role of firm heterogeneity in CZ effects. We then present our approaches to examining city size effects, the role of industry differences in CZ-average wages, and wage differences between more- and less-educated workers.

Two-way fixed effects model

Building on AKM and an extensive body of subsequent work, we start with a two-way fixed effects model of (log) earnings determination with additive employee and employer fixed effects. (Recall that our sample excludes people with multiple employers in a quarter.) Letting y_{it} represent the log of observed earnings of worker i in quarter *t* and letting *f(i,t)* be an index function that gives the identity of the establishment where she is working, our baseline AKM specification is:

$$y_{it} = \alpha_i + \delta_{f(i,t)} + X_{it}\beta + \epsilon_{it}. \tag{1}$$

Here, α_i is a fixed effect that captures the time-invariant skills (measured or unmeasured) of worker i, X_{it} is a vector of time-varying characteristics (age and calendar time effects), and δ_f is a fixed effect

[13] Glaeser and Maré (2001) emphasized the possibility of an adjustment process for movers. As can be seen in Appendix Figure 1, their NLSY sample shows that it takes about a year to get to the new level of earnings for people who enter or leave a metro area.

that captures the wage premium paid by establishment f.[14] The error term ϵ_{it} captures all other factors, including transitory worker-specific earnings shocks; transitory firm-, industry-, or CZ-wide shocks; and any person-specific match effect associated with working in the specific firm (or CZ or industry).

Ordinary least squares (OLS) applied to equation (1) will yield unbiased estimates of the establishment premia if the $\epsilon_{it}'s$ are orthogonal to the sequence of workplaces selected by worker i -- a so-called "exogenous mobility" (EM) assumption. Card, Heining, and Kline (2013) and Card, Cardoso, and Kline (2016) develop a number of specification tests that address the plausibility of this assumption. We reproduce some of these tests in the Appendix, and show that model (1) with exogenous mobility provides a relatively good approximation to the patterns of observed earnings outcomes for movers between workplaces, though it is clearly an over-simplification. We provide a detailed discussion of the patterns of residuals for CZ movers below.

Measuring CZ wage premia

We define the locational wage premium associated with CZ *c* as a weighted average of the establishment effects in that CZ:

$$\Psi_c \equiv \frac{\sum_{c(f)=(c)} N_f \delta_f}{\sum_{c(f)=(c)} N_f} \ , \tag{2}$$

where *c(f)* is an index function giving the CZ for establishment *f* and N_f is the number of person-quarter observations in our estimation sample for that establishment. This premium can be interpreted as follows: If a worker were selected at random from CZ *c* and moved to a randomly selected workplace in CZ *c'*, her earnings would increase, on average, by $\Psi_{c'} - \Psi_c$.

Importantly, by defining Ψ_c in this way, we abstract from problems caused by heterogeneity in workplace pay premiums within CZs. To see how such heterogeneity can matter, decompose the establishment pay premium in the basic AKM model (1) into the sum of the CZ premiums and a residual:

[14] We normalize the establishment effects so that the person-quarter-weighted average of δ_f across all establishments in the large, low-wage restaurant industry (NAICS code 7225) is zero.

$$\delta_f = \Psi_{c(f)} + h_f. \tag{3}$$

We refer to h_f as the *hierarchy* component of earnings – it reflects workplace *f*'s position in the local job ladder. Now consider the change in earnings for a worker who moves across CZs between period t-1 and t. Using equations (1) and (3) (and ignoring any change in the X's) this can be written as:

$$\begin{aligned} y_{it} - y_{i,t-1} &= \left(\delta_{f(i,t)} - \delta_{f(i,t-1)}\right) + (\epsilon_{it} - \epsilon_{it-1}) \\ &= \left(\Psi_{c(f(i,t))} - \Psi_{c(f(i,t-1))}\right) + \left(h_{f(i,t)} - h_{f(i,t-1)}\right) + (\epsilon_{it} - \epsilon_{it-1}). \end{aligned} \tag{4}$$

Notice that the right hand size of (4) includes both the change in the average CZ effect (as defined in equation 2) and the change in the hierarchy effect between the origin and destination workplaces.

In a two-way fixed effects model with person and place effects (rather than establishment effects), the place effects are identified by the mean changes in earnings for movers between places. As shown in equation (4), changes in the hierarchy effects for CZ movers are part of the residual. Only if these changes are orthogonal to the changes in Ψ_c will the mean change in earnings for CZ movers yield unbiased estimates of the place effects for a randomly selected worker. Unfortunately, as we show below, workers who move across CZs typically transit between non-representative firms in their two CZs, such that the change in h_f is negatively correlated with the change in $\Psi_{c(f)}$. We document that failure to account for this leads to attenuation of the estimates of Ψ_c.

<u>Wage Premiums by CZ and Industry</u>

As noted in the introduction, a key issue in urban economics is the role of industrial composition in CZ earnings differences. Our definition of CZ effects means that a CZ with a relatively high share of jobs in high-wage industries (Card et al., 2023) has a relatively high place effect, even if the rate of pay in each industry is the same as in other CZs. To address this, we estimate CZ-by-industry average wage premia,

$$\psi_{cj} \equiv \frac{\sum_{cj(f)=(c,j)} N_f \delta_f}{\sum_{cj(f)=(c,j)} N_f}, \tag{5}$$

where *cj(f)* is an index function giving the CZ (c) and industry (j) of establishment f. We then use these CZ-by-industry wage premiums to explore whether differences in industrial composition can account for differences in Ψ_c. A virtue of our approach to measuring Ψ_c and ψ_{cj} is that the average of ψ_{cj} across all industries in CZ *c* necessarily equals Ψ_c.

Evaluating the Specification

How well do the estimated place effects derived from equations (1) and (2) account for the wage changes of between-CZ movers? To assess this, we begin by dividing CZs into 20 vingtiles based on their estimated Ψ_cs. We then construct the mean changes in the average value of Ψ_c for each of the 20x20 possible origin/destination cells of CZ movers, and relate these to the mean changes in earnings for CZ movers, and the mean changes in the components of equation (4). Because, as noted above, the change in the hierarchy effect for CZ movers tends to be negatively correlated with the change in Ψ_c, we find that the changes in mean earnings for CZ movers are attenuated relative to the mean changes in Ψ_c.[15] However, we show that the changes in hierarchy-adjusted earnings, $y_{it} - h_{f(i,t)}$, are well predicted by the change in the CZ wage effect Ψ_c, suggesting that the estimated place effects are approximately unbiased.

A standard AKM model assumes that only a worker's current establishment (and thus only her current CZ) matters for determining her current wage (di Addario et al., 2023). As noted by De la Roca and Puga (2017), however, is possible that past work experience in certain locations (e.g., the biggest cities) affects wages in all subsequent jobs. Our model absorbs one component of any such dynamic effect, arising through sorting across workplaces. To the extent that experience in big cities provide an entrée into higher-paying firms in other cities, part of the big city experience effect is captured by the

[15] We can further decompose $h_f \equiv \delta_f - \Psi_{c(f)}$ into a component reflecting an industry's position within the CZ and a component reflecting the firm's position within the industry, $h_f = \left(\psi_{cj(f)} - \Psi_{c(f)}\right) + \left(h_f - \psi_{cj(f)}\right)$. We find that changes in both components are negatively correlated with changes in Ψ_c, but that the second component, the firm's location within the CZ-industry pay distribution, is quantitatively more important.

establishment effects and associated hierarchy terms in (1). However, there may also be a dynamic component, reflecting (for example) better match effects for workers with big city experience.

To explore the impacts of allowing for dynamic place-specific experience effects, we estimate a version of model (1) with controls for the number of quarters of previous employment (regardless of where that work was done), and for quarters of previous employment in specific sets of CZs (e.g., the 10 largest or 25 largest CZs), interacted with indicators for where the worker is currently located. Because our LEHD data start in 2010, our work histories are incomplete for all but the youngest workers. Thus, this analysis is restricted to those who were 26 or younger in 2010. While we find evidence of dynamic effects, they are relatively small, and their addition leads to little change in the estimated contemporaneous effects for working in different CZs. Accordingly, for most of our analysis we use the simpler specification without such dynamic effects, fit to our overall sample of workers.

Decomposing Mean Earnings Differences Across CZs

Assuming that the specification of equation (1) is valid, we can use it to decompose differences in mean wages across CZs. Averaging across workers and time periods, we get:

$$\bar{y}_c = \bar{\alpha}_c + \Psi_c + \bar{X}_c\beta \tag{6}$$

where $\bar{\alpha}_c$ – the mean of the person effects for workers in the CZ – summarizes the average skill of the workforce, Ψ_c represents the locational premium (defined in 2), and $\bar{X}_c\beta$ represents a coefficient-weighted average of the time-varying effects (which we expect to be nearly constant across CZs).[16] This decomposition differs from one based on a simple cross-sectional model (such as the one described in Section II) because it uses the fixed effects for each worker to measure both the observed and *unobserved* components of skill. We measure the share of the variation in mean earnings that is

[16] Note that because establishments are nested within CZs the CZ averages of the hierarchy and residual terms are mechanically zero.

attributable to each of the components on the right-hand side of (6) by regressing each component on CZ mean earnings.

We can also use (6) to decompose the overall variation in CZ mean wages. Omitting terms that are mechanically zero, the across-CZ variance of mean earnings is:

$$\begin{gathered} V(\bar{y}_c) = V(\bar{\alpha}_c) + V(\Psi_c) + V(\bar{X}_c\beta) \\ +2cov(\bar{\alpha}_c, \Psi_c) + 2cov(\bar{\alpha}_c, \bar{X}_c\beta) + 2cov(\Psi_c, \bar{X}_c\beta) \end{gathered} \tag{7}$$

Of particular interest are the relative magnitudes of $V(\bar{\alpha}_c)$, $V(\Psi_c)$, and $cov(\bar{\alpha}_c, \Psi_c)$), the last of which measures the tendency of higher-skilled workers to live in CZs that pay higher wage premiums. [17]

Decomposing city and industry effects

Next we consider the structure of the estimated CZ wage effects, focusing on three related questions: (i) Is the CZ premium Ψ_c constant across industries? (ii) Is the average wage premium for certain cities partly due a higher share of jobs in industries that tend to pay more everywhere (e.g., finance in NYC)? (iii) To the extent that ψ_{cj} varies across CZs, is this variation associated with the share of local employment in the sector (i.e., a local concentration effect)?

We begin with a simple analysis of variance of the estimated $\psi_{cj}s$ from (5):

$$\hat{\psi}_{cj} = \theta_c + \mu_j + e_{cj} \,. \tag{8}$$

The R-squared of this model provides an initial indication of the importance of CZ-industry match effects, which are captured in the residual term e_{cj} (along with any sampling error in $\hat{\psi}_{cj}$).

We also estimate specifications that include the share of employment in industry *j* in CZ *c* (or the log of this share) as an additional regressor in (8). The coefficient on this variable measures the extent to which the excess pay premium for a given industry in a given CZ (i.e., the match effect) is related to the local size of an industry, reflecting potential agglomeration effects.

[17] Variance decompositions like (7) are biased when estimated at the level of the individual worker (Andrews et al., 2008; Kline et al., 2020). These biases derive from the small samples used to estimate each individual establishment and worker effect, and should not be present in decompositions of CZ aggregates. We have verified this empirically for a somewhat different specification in an earlier version of this paper (Card et al. 2022).

To quantify the effects of industry composition and CZ-specific industry pay premiums on the variation in average CZ pay effects, we turn to an Oaxaca (1973) style decomposition. Let w_c represent the share of national employment in CZ c; let s_{cj} and $\bar{s}_j$ represent the shares of CZ and national employment, respectively, in industry *j*; and let $\bar{\psi}_j$ represent the national average wage premium for industry *j*, $\bar{\psi}_j \equiv \sum_c w_c \psi_{cj}$.[18] The city average wage premium can then be decomposed as:

$$\Psi_c = k + \underbrace{\textstyle\sum_j \bar{s}_j (\psi_{cj} - \bar{\psi}_j)}_{1} + \underbrace{\textstyle\sum_j (s_{cj} - \bar{s}_j)\bar{\psi}_j}_{2} + \underbrace{\textstyle\sum_j (s_{cj} - \bar{s}_j)(\psi_{cj} - \bar{\psi}_j)}_{3}. \tag{9}$$

where $k \equiv \sum_j \bar{s}_j \bar{\psi}_j$ represents the average industry premium earned by workers in a representative CZ with national-average industry shares, and is constant across CZs.

Term 1 on the right-hand side of (9) is a share-weighted average of the gap between the industry premium earned in CZ *c* and the corresponding national average premium, $(\psi_{cj} - \bar{\psi}_j)$. This measures the <u>average</u> excess pay premium for workers in each industry, relative to the national premium in the corresponding sector, and can be interpreted as a "composition adjusted" local pay premium – the gain or loss in wages that a typical worker would see if he or she moved from another CZ to a representative firm *in the same industry* in this CZ. Term 2 represents the excess share of workers in different industries $(s_{cj} - \bar{s}_j)$, weighted by the national average wage premium in each industry. This is a measure of the industry composition effect in Ψ_c, attributable to higher or lower shares of employment in high-wage or low-wage industries. Finally, term 3 arises from a correlation between the share of an industry in a given CZ and any excess local wage premium in that industry. This correlation will be positive if there are higher CZ-specific premia for industries with higher local employment shares – as suggested by models of agglomeration externalities. We refer to these three terms below as the *average earnings premium*, the *composition effect*, and the *interaction effect*.

[18] Note that we use a fixed set of city weights to define the national average wage premiums for each industry, rather than using industry-specific city weights. This choice simplifies the interpretation of the interaction term in our decomposition.

To illustrate, consider the example of New York City, which has an above-average share of employment in the finance industry. Term 1 will be larger if finance *and other sectors* pay more in NYC than do jobs in the same sectors in the average CZ. Term 2 will be larger if NYC has a relatively large share of workers in finance and other high-paying industries. Finally, term 3 will be larger if industry-specific pay premiums in NYC are higher, relative to other cities, in the industries (like finance) where NYC workers are more likely to be employed.

To summarize the implications of equation (9) we calculate the variances and covariances of the three non-constant terms on the right-hand side and compute the shares of each term in the overall variance of Ψ_c. One of the covariances is of special interest: that between terms 1 and 2. Some models (e.g., Beaudry et at., 2012) suggest that an increased share of employment in high-wage industries will raise wages in all industries (not just in the sector itself, as is predicted by agglomeration theories). Such a spillover effect from the presence of high-wage industries will appear as a positive covariance between term 2 and term 1.

Decomposing the effects of CZ characteristics on average earnings

The framework of equations (6) and (9) provide a useful way to understand the channels through which CZ characteristics affect average earnings in the CZ. For example, a long tradition in urban economics explores the tendency for average wages to rise with city size (e.g., Behrens et al. 2014; Eeckhout et al. 2014; Butts et al. 2023). Equation (6) implies that if we fit a simple regression model like

$$\bar{y}_c = \delta_0 + \delta_1 Z_c + \xi_c, \tag{10}$$

relating mean log earnings in a CZ to a characteristic Z_c like the log of CZ population, then the overall effect of Z_c can be decomposed into the sum of the effects of Z_c on $\bar{\alpha}_c$, Ψ_c, and $\bar{X}_c\beta$. Moreover, equation (9) implies that the effect of Z_c on Ψ_c can be further decomposed into an effect on the composition adjusted local pay premium (term 1 in equation 9), an effect on the industry composition component of the CZ pay premium (term 2) and an effect on the interaction effect that captures potential agglomeration mechanisms (term 3). We explore decompositions like this for the size gradient as well as the gradients of CZ premiums with respect to mean earnings and mean education.

Geographic variation in the return to education

In our final step, we examine differences in the returns to education across places. The starting point for this analysis is the fact, noted in Section II, that the gap in wages between more- and less-educated workers is bigger in larger cities and in cities with higher average wages. Is this because the local pay premiums are actually different for higher- and lower-educated workers? Or is it due to differences in the unobserved skills of more- and less-educated workers in different cities?

To assess this, we use the AKM estimates of firm effects from (1), but construct education-group specific weighted averages of the firm effects in each CZ and CZ-industry pair. [19] Denoting education by e (e=H, L), the CZ effect for education group e is:

$$\Psi_{ce} = \frac{\sum_{c(f)=(c)} N_{fe}\delta_f}{\sum_{c(f)=(c)} N_{fe}}, \tag{11}$$

where N_{fe} represents the number of workers of education group e at establishment f.

Using equation (1), the average earnings gap across education groups in CZ c, $\bar{y}_{cH} - \bar{y}_{cL}$, can be written as:

$$\bar{y}_{cH} - \bar{y}_{cL} = (\bar{\alpha}_{cH} - \bar{\alpha}_{cL}) + (\bar{X}_{cH} - \bar{X}_{cL})\beta + (\Psi_{cH} - \Psi_{cL}) + (\bar{\epsilon}_{cH} - \bar{\epsilon}_{cL}). \tag{12}$$

The first term in (12) reflects the gap in average human capital between high- and low-education workers in the CZ, the second term reflects differences in the covariate effects for the two groups, and the third reflects difference in the average pay premiums for the two groups (i.e., in the premiums

[19] We use education data collected in the ACS and linked to LEHD, which are available for about 10% of all workers in LEHD. See Abowd et al. (2009). This small coverage prevents us from estimating (1) separately by education. Instead, we use the full sample for estimation of the AKM model, then average δ_f separately weighting by the number of high- and low-education workers observed at each firm. We explore the restriction of common establishment premiums across education in the Appendix.

offered by the firms where the two groups are employed). The final term reflects any change in residuals from before to after a move between CZs.[20] We expect it to be small.

Generalizing equation (9), we decompose the third term in (12) as:

$$\Psi_{cH} - \Psi_{cL} = \underbrace{\sum_j s_{cj}(\psi_{cjH} - \psi_{cjL})}_{1} + \underbrace{\sum_j (s_{cjH} - s_{cjL})\psi_{cj}}_{2}$$
$$+ \underbrace{\sum_j (s_{cjH} - s_{cj})(\psi_{cjH} - \psi_{cj}) - (s_{cjL} - s_{cj})(\psi_{cjL} - \psi_{cj})}_{3}. \qquad (13)$$

Here, ψ_{cj} is the CZ-industry premium as defined in (4), ψ_{cje} is the education-group-specific CZ-industry premium defined in (11), and s_{cje} represents the share of group-*e* workers in CZ c who work in industry j (with $\sum_j s_{cje} = 1$ for each c and e). The first term in (13) reflects an average "education premium" in the CZ (i.e., a weighted average of the difference in pay premiums for high and low educated workers in the same industry in that CZ). The second term is an industry composition effect, reflecting the relative shares of the two groups in higher premium industries. The third is a variant of the interaction effect in equation (9) with components reflecting the *relative* clustering of high- versus low-education workers in industries with a higher or lower local relative industry premia for that group.

Combining (12) and (13) allows us to decompose the mean pay gap between more and less educated workers in a CZ into six components, two of which (the $X\beta$ and $\bar{\epsilon}$ terms) we expect to be quite small. We use this for two purposes. First, we explore which of these components account for the association of the local return to education with the CZ size or average wage level by regressing each component separately on CZ size or mean wage. Second, we decompose the across-CZ variance in $\bar{y}_{cH} - \bar{y}_{cL}$ into the variances of each of the components (plus covariances).

[20] Recall that the mean of ϵ_{it} within workers is zero, so non-movers do not contribute to this term; similarly, hierarchy terms average to zero within CZs

V. CZ Earnings Premiums

This section introduces our estimated CZ earnings premiums and presents a range of evidence supporting their validity. In Section VI, we discuss the role of industry composition in CZ earnings differences, while Section VII discusses variation in the CZ return to education.

A. The AKM decomposition and the geographic structure of establishment earnings premiums

We estimate equation (1) using our main LEHD sample, which has 2.5 billion person-quarter observations for about 110 million different workers. Table 2 presents a variance decomposition of the terms of (1), first at the original person-quarter level in columns 1 and 2 and then aggregated to the CZ mean level in columns 3 and 4. At the person-quarter level, the standard deviation of log quarterly earnings is 0.653. About 64% of the variance is attributable to the person effects, 12% to firm effects, and 4% to the time-varying covariates (age and time effects). Another 10% is attributed to the positive covariance between the person effects and the firm effects, while the covariances involving $X_{it}\hat{\beta}$ are small, as expected. Finally, 13% of the overall variance is unexplained (implying an R-squared coefficient for the model of about 87%). These variance shares are not too different from those that have been estimated in other settings (e.g., in data from West Germany in the 2002-2009 period, presented by CHK). In particular, the positive correlation between the employer and employee pay components we see in the U.S. data is similar in magnitude to that in recent data from West Germany, Portugal, and Italy (see Card et al., 2018).

When we aggregate the data to the CZ level, in columns 3-4, the variance share accounted for by person effects falls, but worker skills still account for 30% of the variance of mean earnings across CZs. Firm effects (which, when averaged to the CZ, form our estimates of CZ effects) explain nearly the same share, much more than so than at the individual level. An implication is that there is substantial geographic variation in average establishment premiums. Indeed, about 13% of the variance of establishment effects is between CZs. A separate analysis, not reported in the table, indicates that about 23% is between two-digit industries and that about 40% is between CZ-by-industry cells.

Importantly, nearly 40% of the variance of CZ mean earnings is attributable to the positive covariance between person effects and CZ effects.[21] The correlation between $\bar{\alpha}_c$ and Ψ_c is 0.64, much higher than the corresponding correlation (0.18) at the individual level, implying that there is relatively more sorting of high skilled workers to high-paying places than there is sorting of high skilled workers to high-paying workplaces within CZs.[22]

An alternative to the variance decomposition in columns 3 and 4 of Table 2 is a covariance decomposition based on equation (6): $var[\bar{y}_c] = cov[\bar{y}_c, \bar{\alpha}_c] + cov\left[\bar{y}_c, \widehat{\Psi}_c\right] + cov[\bar{y}_c, \bar{X}_c\beta]$. The relative contributions of the covariance terms are just the regression coefficients from (weighted) regressions of the mean person effects and CZ average wage effects on CZ average wages. (The third covariance term is approximately 0.) We present these coefficients in the second and third rows of column 1 in Table 3. Their values, 0.499 and 0.489, respectively, indicate that approximately equal shares of the variation in mean CZ earnings are attributable to differences in worker skill characteristics captured in $\bar{\alpha}_c$ and to differences in local wage premiums, $\widehat{\Psi}_c$. As noted in Section II, a similar exercise based on a richly specified cross-sectional model fit to the ACS shows that <u>observed</u> skills can account for only 31% of the variation in mean CZ wages. Clearly, models that ignore unobserved skills can significantly understate the degree of skill-based sorting across places.

---

[21] In a similar decomposition of labor market mean wages in West Germany, Dauth et al. (2022) report that the variance shares explained by person effects, establishment effects, and the covariance of person and establishment effects are 39.8%, 23.6% and 41.7%, respectively – suggesting that the contribution of worker skill composition to differences in mean wages across places is slightly larger in Germany than the U.S.

[22] A caveat is that the individual level correlation between the estimated employee and employer fixed effects is downward biased by estimation error (see Kline et al, 2022), whereas at the CZ level this bias is negligible. However, evidence in Gerrard et al. (2022) suggests that the degree of attenuation bias in the worker-level correlation is modest with reasonable numbers of wage observations per person (as we have in our sample).

The second and third columns of Table 3 present regressions of $\bar{y}_c$, $\bar{\alpha}_c$, and $\widehat{\Psi}_c$ on CZ size (column 2) and the share of college-educated workers in the CZ (column 3). These yield a similar pattern: Half or more of the gradient of log earnings with respect to CZ size or the CZ college share is attributable to differences in person effects, rather than to place effects. This is much larger than the 20% share of the CZ size gradient that we found attributable to observed worker skill characteristics in the ACS in Section II. The degree of worker sorting across CZs is substantially under-estimated by focusing only on observed skills.

The lower rows of Table 3 show that higher-wage, larger, and higher-education CZs have more dispersion of worker effects, with many more workers in the top national decile and somewhat fewer in the bottom decile. These CZs also have stronger assortative matching of high-skill workers to high-premium firms (see the bottom row of the table). The relationship of $correl(\alpha_i, \delta_{f(i,t)})$ with the college share is particularly impressive: there is stronger assortative matching in high-education places.

To help visualize our estimates of CZ earnings differentials, Figure 2 presents two simple maps. The upper panel classifies CZs into five quintiles based on mean log wages in our ACS sample. The lower panel shows the classification of the same CZs into quintiles of the estimated average CZ premium – i.e., $\widehat{\Psi}_c$. Table 4 shows the characteristics of CZs in each quintile, as well as those with very high (top 10) and very low (bottom 10) estimated CZ effects.[23] We also identify two subgroups of top-10 wage premium CZs: larger urban areas (like New York or San Francisco), characterized by high place premiums and relatively highly skilled workforces; and resource-intensive CZs (like oil-producing areas in Texas) where workers have just average skills.

[23] For disclosure reasons we cannot report individual CZ averages, and the classification in columns 2-5 is limited to CZs with at least 25,000 people. The top ten CZs consist of four large urban CZs (Arlington, VA; New York, NY; San Francisco, CA; and San Jose, CA) and six that we classify as "resource intensive" (Anchorage, AK; Elko, NV; Hobbs, NM; Odessa, TX; Pampa, TX; and Rock Springs, WY). The bottom ten CZs are Boone, NC; Campbellsville, KY; Ocala, FL; Statesboro, GA; Presque Isle, ME; West Plains, MO; Poplar Bluff, MO; Eldon, MO; Missoula, MT; and St. George, UT.

The maps illustrate the positive but imperfect correlation between CZ premiums and mean log wages. This reflects both the direct effect of the pay premiums and an indirect effect arising from assortative matching of highly skilled workers to higher-premium CZs. The estimates in Table 2 imply that a 1 percentage point increase in a CZ's average pay premium is associated with a 0.65 percentage point increase in earnings through the presence of higher-ability workers, yielding a total wage increase of 1.65 percentage points. There are, however, some interesting exceptions to this pattern. Florida has mostly middle or upper quintile average wages but bottom quintile CZ effects, while many resource-intensive areas (e.g., the agricultural regions of California, or the Texas oilpatch) have middle or low average wages but top quintile CZ premiums.

Census disclosure rules prevent us from reporting exact values of the CZ-specific estimates from the LEHD. As an alternative, we have constructed predictions of the LEHD-based premiums from publicly available ACS data. Our procedure is described in the Appendix. The predictions are correlated 0.90 with the LEHD premiums, and are scaled as unbiased predictions. Appendix Table 3 reports the predicted premiums for the 50 largest CZs. It shows that coastal CZs like New York, San Francisco, San Jose, and Washington DC have premiums of 10-18% relative to the average CZ, while Cleveland, Pittsburgh, Buffalo, and several Florida CZs pay 4-6% less than the average CZ. A data file with predictions for all 691 CZs identifiable in ACS data will be distributed with this paper.

B. Validating the Specification

The results in Tables 2 and 3 all derive from our AKM specification (1), and so are predicated on its underlying assumptions. As noted in Section IV, for OLS to provide unbiased estimates of the CZ wage premiums we need exogenous mobility (EM) to hold: Moves across CZs or firms have to be uncorrelated with the error term ϵ_{it} in (1). In this section we examine earnings changes around CZ moves for evidence of violations of EM. The appendix presents extensive analyses of firm-level moves. To foreshadow our results, we find that EM can be rejected, but that departures from the patterns predicted by EM tend to be small.

Event studies

Figure 3 presents average earnings by quarter relative to the date of a move, separately for those who move across different quartile groups of CZs, grouped by Ψ_c. It parallels Figure 1, but categorizes CZs based on their estimated wage premiums rather than on mean earnings. As in Figure 1, we show only the earnings profiles for movers from the top and bottom quartiles.

We see again that earnings are flat prior to a move and but adjust a bit more slowly following a move, increasing for several quarters. We also see that earnings changes from period -1 to period 1 are consistent in direction and relative magnitude with the estimated CZ premiums: Movers from quartile 1 to 4 see substantial earnings increases, while movers from 4 to 1 see substantial declines. There are similar but smaller changes for less extreme moves – e.g., those who move from quartile 4 to 3 see small earnings declines, while movers from 4 to 2 are intermediate.

The large scale of Figure 3 makes it difficult to assess the exact fit of the model for different groups. Appendix Figure 6 repeats the exercise, this time showing mean residuals from model (1). It shows that all groups see sharp declines in their residuals in quarter +1 that largely recover by quarter 2, consistent with a temporary negative effect of mobility on earnings. There is also a tendency for those who move down the CZ premium distribution to see small increases in their residuals following the move, relative to those who move in the opposite direction, consistent with Roy-style sorting. But these differentials are small in magnitude – for example, the residuals of 4-to-1 movers rise by about 1.5 log points relative to those of 1-to-4 movers, one-tenth or less of the magnitude of the difference in CZ premia.

Dynamic effects

One possible explanation for this pattern of residuals around a move is that movers from upper quartile CZs receive a return to their experience in these places. To explore this possibility, we estimated a set of models that included four additional control variables: (i) cumulative quarters of work for each individual (as measured from the start of our LEHD panel); (ii) cumulative quarters of work in larger CZs; (iii) an interaction of cumulative work experience with an indicator for currently working in a larger CZ; (iv) an interaction of cumulative work experience in larger CZs with an indicator for currently working in a larger CZ. These models generalize De la Roca and Puga's (2017) specification slightly by allowing the returns to overall experience and "big city" experience to be valued differently in larger and smaller CZs.

Unfortunately, our data do not allow us to identify workers' locations or work histories prior to 2010. Thus, for this analysis we limit attention to workers under age 26 in 2010, for whom our experience measures are close to complete. We begin by re-estimating our static model (equation (1)) on this subsample and re-computing the implied CZ premia for younger workers. The first column of Table 5 shows that the standard deviation of CZ premia computed from the young-worker sample is 0.078, very similar to the 0.079 for the full sample (reported in Table 2). A regression of the young-worker estimates of Ψ_c on those from the full sample has a coefficient 0.96 and R-squared of 0.96 (shown in the second and third rows), indicating that the change in sample makes very little difference to our estimation of CZ premiums.

Columns 2 and 3 report specifications augmented with the dynamic experience controls. We test two definitions of "larger CZs": the 10 largest CZs (column 2); and the 25 largest CZs (column 3).[24] The estimated experience coefficients from the dynamic models are presented in the lower part of the

[24] The 10 largest (in order) are Los Angeles; New York; Chicago; Washington DC; Northern NJ; Houston; Philadelphia; Boston; San Francisco; and Atlanta. They account for 24% of the U.S. population. The CZs around the top 25 cutoff are San Jose (#23), Cleveland (#24), St. Louis (#25), Pittsburgh (#26), and New Orleans (#27). The top 25 CZs account for 40% of the U.S. population. De la Roca and Puga's (2017) analysis of Spanish labor market data also considers two groups of the largest cities. Their first group consists of Madrid and Barcelona (26.3% of the population); their second adds the next three largest cities (for a total of 37.9%).

table. (We show estimated standard errors for these coefficients that are clustered by the first CZ in which an individual is observed in our sample). These coefficients point to three main conclusions. First, focusing on earnings of workers outside of the larger CZs, cumulative work experience has a relatively strong, statistically significant effect on earnings, with an implied return of around 1.6 log points per year of work (on top of the returns to age captured in our main covariate index).[25] Second, experience in a larger CZ has about double that effect, totaling about 3.3 log points per year of work, for workers who are currently working outside larger CZs. Third, whether big city work experience is differentially valued in larger or smaller CZs is depends on the definition of "large CZs". The specification in column 2 indicates that large-CZ experience has a *lower* return for workers who remain in those CZs, but this is not the case for the specification in column 3.

As in column 1, the estimates in the upper rows of columns 2 and 3 summarize regressions of the estimates of $\widehat{\Psi}_c$ from the dynamic specifications on the place effects from our preferred model and sample. We find that the addition of the dynamic experience controls makes essentially no difference – the slope coefficients, R-squared, and standard deviations are nearly identical to those in column 1, which are based on the same young-worker sample but omit dynamic controls. We conclude that there is very little systematic bias from using the $\widehat{\Psi}_c$ estimates from our main specification. We thus focus on these estimates, and on the full sample, in the remainder of this paper.

A closer look at earnings changes for CZ movers

Another useful check on the CZ wage premiums we obtain from our AKM model is to look at short-run earnings changes around moves, and compare them to the predictions from the model. When a worker moves from CZ c to c', do her earnings immediately rise by the full amount $\Psi_{c'} - \Psi_c$? We already noted in our discussion of Figure 3 that the first quarter following a move appears to show

[25] Because we also control for both age and calendar quarter, the experience main effects are identified by variation in the cumulative number of quarters of employment over the period covered by our sample, and so reflect in part consistency of labor force attachment.

anomalously low earnings, so this is a stringent test. Nevertheless, we use our event study sample to compare the observed earnings changes of CZ movers to the predictions based on our estimated CZ effects. For disclosure reasons we cannot report the Ψ_cs for individual CZs, so instead we group CZs into vingtiles based on their Ψ_cs and compare the 400 possible origin-destination combinations.

Panel A of Figure 4 plots age-adjusted earnings changes (i.e., the changes in $y_i - \widehat{X_\iota}\hat{\beta}$) for CZ changers against the origin-to-destination mean changes in the CZ effects for each of the 400 cells. We overlay a 45-degree line, which would correspond to earnings changes that exactly matched the change in CZ premiums. The scatterplot is much flatter than this, with a slope of just 0.56: On average, CZ movers see just over half of the change in earnings, upward or downward, that we would have predicted from the change in CZ premia.[26]

Equation (4) shows that age-adjusted changes in earnings for CZ movers can be decomposed into three components: The change in the CZ premium, the change in the hierarchy effect, and the change in the AKM residual. By construction the first component changes one-for-one with the change in CZ premium for movers. Panel B examines the third component, the residual. Evidence that the change in residuals varies systematically with the change in CZ effects would indicate violations of the EM assumption. We see that the scatter of points is quite flat, with a slope of just -0.06. Nevertheless, the slope is statistically significant, pointing to somewhat less than full adjustment of earnings when workers move to higher- or lower-wage CZs. Thus, as with the other evidence above and in the appendix, we infer that the EM assumption is not exactly right, but is reasonably accurate.

Panel C plots the change in the hierarchy component against the change in the CZ effect for each of the 400 origin-destination groups. This has a much stronger negative slope of -0.38. Workers

[26] This is not inconsistent with the pattern in Figure 3: Although movers to higher-premium CZs see earnings increases (the slope in Figure 4, panel A is positive), these changes are smaller than would be predicted from the change in CZ premiums. Note that the scatterplot in Figure 4, panel A is quite tight, suggesting that the *ranking* of earnings changes, though not the magnitude, is closely approximated by the change in premiums.

who move to higher-premium CZs tend to land at firms that are lower-ranked in the local wage ladder than are the firms from which they departed (i.e., their hierarchy effect becomes more negative), while workers who move to lower-premium CZs experience gains in their relative hierarchy effect.

What can explain this pattern? Card et al. (2023) propose a simple model in which potential movers evaluate job offers based on the expected wage they will receive on a new job relative to their old job. In essence, job searchers who only care about their total wage will be willing to trade off different components of wages, leading to a negative correlation between the changes in one component and the changes in any other. Since the total wage includes the CZ place effect for the job, plus the hierarchy effect for the workplace within the CZ, plus any expected match effect, such a model could potentially explain both the negative correlation between the change in hierarchy effects and the change in CZ effects in Panel C of Figure 4, and the negative correlation between the change in the AKM residual and the change in CZ effects in Panel B.

Based on the patterns in Figure 4, it seems that job changers are more willing or able to substitute between place effects and hierarchy effects than they are to trade off place effects against the AKM residual component of wages. A further decomposition of the hierarchy component into a local industry effect and the deviation of the workplace wage premium from the average for its industry shows that -0.30 of the -0.38 slope (~80%) of the change in hierarchy effects shown in Panel C of Figure 4 reflects a change in the idiosyncratic workplace pay premium within industry and CZ; only -0.08 (~20%) is due to changes in the average industry premium in the destination CZ relative to the origin CZ: Workers trade off CZ premiums against within-industry firm premiums more than they do against industry average premiums.

Panel D of Figure 4 shows results for the hierarchy-adjusted change in earnings of CZ movers. These increase almost one-for-one (slope = 0.94) with CZ premiums. This is the basis for our conclusion that violations of the EM assumption are quantitatively small relative to CZ premiums. Whatever the failings of the AKM specification, it does a good job of capturing CZ-level wage premiums, once we take account of heterogeneity in the pay premiums offered by different employers in the origin and destination labor markets.

C. Comparisons of CZ wage premiums from alternative approaches

Most existing studies of place-based earnings differentials use either cross sectional data (such as the ACS data used in Section II, above), or estimate two-way fixed effects models that include person effects and place effects. In particular, Combes et al. (2008), De la Roca and Puga (2017), and an earlier version of this paper (Card et al. 2022) all use models with metro area effects.[27] The pattern of changes in hierarchy effects for CZ movers shown in Figure 4 suggests that such models will lead to attenuated place effects because the change in hierarchy effects for CZ movers, which are treated as part of the error term in specifications with person and place effects, is strongly negatively correlated with the change in place effects. We verify this in Table 6, comparing our AKM-based estimates of Ψ_c to estimates from specifications that regress log wages on individual and CZ (column 2) or CZ-by-industry (column 3) effects. These results confirm that "person and place effect" models yield CZ average pay premiums with significantly smaller standard deviations than our establishment-based approach. Moreover, when we regress the estimates of Ψ_c from these specifications on the estimates from our preferred approach, the coefficients are 0.77 and 0.79, respectively, suggesting that the gap in premiums between high-wage and low-wage CZs is understated by 20%. Nevertheless, the R-squared statistics from these models are reasonably high, suggesting that the ranking of CZs is fairly similar in the alternative and preferred specifications.

It is also interesting to see how the place effects from simple cross-sectional models are related to the estimates of Ψ_c from our preferred approach. Since we have few covariates in the LEHD, we use

[27] The seminal study by Glaeser and Mare (2001) only distinguished between metro and non-metro areas. To the best of our knowledge, only Dauth et al. (2022) use an AKM-style firm-based approach to estimate place effects.

the ACS sample described in Section II, above, to estimate three simple models: one with basic human capital controls (a flexible combination of education, gender, race/ethnicity and experience); a second that adds information on immigrant source region and years in the U.S. as well as field of study for people with a B.A. or higher education; and a third that also controls for industry. (The third model is the one highlighted in Section II.) The standard deviations of the estimated CZ effects, and the regression coefficients from models that project the ACS-based place effects on our preferred estimates, are shown in columns 5-7 of Table 6. For reference we also show parallel statistics for the unadjusted mean of log wages in each CZ in column 4.

The results show that even richly specified cross-sectional models, such as the one used in column 7, over-state the variation in place effects in the U.S. – presumably reflecting the role of unobserved ability. Interestingly, however, the degree of "upward bias" in the cross-sectional estimates is not too different from the degree of "downward bias" in specifications that use person and place effects and ignore heterogeneity in workplace pay premiums. Cross-sectional models overstate the variation in place effects, while estimates from person and place effects models understate the variation. As before, *rankings* of cities by cross-sectional estimates are similar to those by our preferred model, reflecting the strong correlation between worker ability ($\bar{\alpha}_c$) and CZ premiums discussed earlier.

VI. The role of industry

Next we explore the role of industry in the estimated ψ_{cj} effects obtained by averaging the establishment pay premium by CZ and industry (see equation 5). Of interest are three issues: the extent to which industry and place effects are additively separable; the role of industry composition in explaining locational wage differentials; and the contribution of industry agglomeration and/or local specialization effects to CZ wage differences.

We begin by using equation (8) to explore potential interactions between industry and place effects. We use a 2-digit industry classification; below, we explore more detailed industries. Table 7 reports models for the estimated $\psi_{cj}'s$ that include CZ effects (column 1), industry effects (column 2), and both set of effects (column 3). The R-squared statistics show that 33% of the variation in $\hat{\psi}_{cj}$ can be explained by CZ, while 57% can be explained by industry. Remarkably, the combined explanatory power of the two sets of effects is 87.3%, just 2.3 percentage points less than the sum of that explained in columns 1 and 2. The implication is that place and industry factors are nearly orthogonal, so the CZ-wide component of ψ_{cj} is similar whether we control for industry or not – a fact that anticipates the relatively small role of industry composition in CZ place effects identified by our more complete decomposition models below. Nevertheless, the less than perfect fit of the model in column 3 means that industry pay premiums vary across CZs. (We expect sampling error in $\hat{\psi}_{cj}$ to be quite small, and in earlier work estimated that the reliability of the match effects in (8) is over 90%.)

One potentially important source of CZ-specific returns to industry – emphasized by Marshall (1920) -- is the relative size of the local industry, which may drive agglomeration effects. To explore this, we add to the model from column 3 a control for s_{cj} – industry *j's* share of employment in the CZ – or for its log.[28] Column 4 indicates that a one percentage point increase in an industry's local employment share is associated with about a 0.45 percent increase in relative earnings in that industry, on top of the nationwide industry wage premium, while column 5 indicates that a one percent increase in industry employment in a CZ is associated with a 0.044 percent increase in relative earnings.[29] Both are significant, and are consistent with the presence of local agglomeration effects. However, the goodness of fit increases only slightly relative to the model in column 3. Local concentration differences explain only about 1-2% of the overall variation in CZ-industry effects. We conclude that although local

[28] Because the model includes CZ fixed effects, controlling for the log employment share is equivalent to controlling for the log number of workers in the industry in the CZ.
[29] These approximately coincide for an industry with 10 percent of employment in the CZ at baseline.

agglomeration is associated with higher earnings, the net effect is small because industry shares do not vary enough across CZs.[30]

To explore this conclusion further, in Table 8 we turn to the formal decomposition of CZ-mean effects given by equation (9). We present both two- and four-digit industry classifications here, with generally similar results. As shown in the second row of the table, pure locational wage premia (term 1 in equation 9) account for over 90% of the variation in Ψ_c. Industry composition effects (term 2 in equation 9) explain only 2.5%, though another 1.5-6.4% is explained by covariance between industry composition and the locational wage premium (i.e., by a tendency for high-wage industries to concentrate in locations that pay high wages across all industries). Finally, specialization of CZs in industries where the CZ has a *comparative* pay advantage (captured by the interaction effects in term 3 of equation 9) explains only 2-4% of the overall variation in Ψ_c, consistent with the modest gain in R-squared noted in the specifications in Table 7 that control for the share or log share of employment in the industry.

Overall, the results in Tables 7 and 8 suggest that the pay premiums for different CZs vary slightly by industry (or equivalently, that premiums for different industries vary a little by CZ), but that differences in industry composition and in CZ-specific returns to different industries are relatively minor factors in understanding the variation in average CZ pay premiums.

Of course there are some potential concerns that could be raised about our approach and the limitations of our LEHD data. One is that we may be incorrectly assigning workers to CZs and/or industries because of establishment imputation errors in the LEHD data. We can assess this by limiting our analysis to observations where we are confident of the CZ and industry assignments. Results of our

[30] We have also explored specifications like those in columns 4 and 5 that allow the industry share coefficient to differ for tradeable vs. non-tradeable industries. The coefficient is slightly larger for tradeable industries, but not substantially so, and there is no meaningful increase in the goodness of fit.

analyses of this subsample are nearly identical to those from our main specification. (See Appendix Table 4 for the alternative version of Table 2.)

A second concern is that our main specification, which ignores dynamic effects of work experience in larger CZs, leads to biases in the implied contribution of local industry structure. To address this we used the estimated CZ-industry effects from the models in Table 5 to redo the decomposition of equation (9). Again, we found results that are very close to our main specification.

A third concern is that our LEHD sample may over-weight individuals who are observed moving in our sample window. We investigated this concern by evaluating the components of the variance decompositions in Table 2 separately for people who are observed in the same CZ throughout our sample. The results, reported in Appendix Table 4, are quite similar for those for our overall sample, leading us to believe that alternative sampling schemes would yield similar results.

A fourth potential concern is that we do not account for occupational differences in our analysis. Although the LEHD lacks information on occupations, we don't believe that such controls are necessary or even desirable. Many occupations (like nurse or attorney) reflect the skills and training that a worker brings to the job. Such characteristics are absorbed in the person effects in our model. Many occupations are also closely associated with specific industries (e.g., auto mechanics in the car repair sector). Our analysis in Table 8 shows that our main conclusions are highly robust to using detailed (4 digit) industry effects to define the pay premiums in equation (9). Finally, insofar as workers gain or lose access to *specific* occupations when they change CZs (such as when a teacher's aide moves to a place with lower credentialing requirements and obtains a position as a classroom teacher) the consequent wage changes could be reasonably interpreted as part of the causal place effect.

To summarize: we interpret the evidence in Tables 7 and 8 as implying that 90-95% of the variation in CZ-specific average wage premiums is attributable to place effects that are common across industries, rather than to the confounding effects of industry structure. Consistent with local agglomeration theories, industry wage premiums are higher in places where an industry employs a higher fraction of local workers, but on net these effects contribute only a few percent of the overall cross-CZ variation in average wage premiums.

This relatively small role for agglomeration may seem inconsistent with some existing studies, such as Greenstone, Hornbeck, and Moretti (2010) (hereafter, GHM). GHM find that the opening of a "million dollar" manufacturing plant leads to a 13% increase in the number of manufacturing plants in the county and to a 6% increase in employment at incumbent plants, for an approximately 19% total increase in manufacturing employment. They also find that a plant opening leads to a 2.7% increase in manufacturing wages (with a standard error of 1.4%). Our estimates (from column 5 of Table 7) imply that a 19% increase in employment would lead to a 0.8% rise in manufacturing wages relative to other industries in the CZ, fully adjusted for the quality of workers. Thus, while our estimate is smaller than GHM's *point estimate* of the total wage effect of a new plant, it is well within the confidence interval for their estimate.

As a final exercise, Table 9 explores the role of industry composition in mediating the effect of CZ characteristics on the mean log earnings in a CZ, examined in Table 3. In the first three rows we reproduce the coefficients from the first three rows of Table 3. These show that 40-45% of the relation between mean log wages in a CZ and the CZ's size or college share is attributable to the place effect in earnings. In the lower rows we then decompose the CZ premium component into the three terms in equation (9). In each case, we find that nearly the entire gradient of Ψ_c with respect to the explanatory variable is due to the composition-adjusted place effect identified in term 1 of equation (9). Industrial composition plays a positive role only in the gradient with respect to CZ mean earnings, and even here it accounts for less than 5%. This further reinforces our conclusion that industry composition differences account for a relatively small part of the differences in average CZ earnings premiums.

VII. Local differences in returns to education

We turn next to an exploration of differences in the return to education across commuting zones. Recall from our discussion in Section II that in the ACS data, the CZ size gradient in raw earnings is steeper for college graduates than for non-college workers; this can alternatively be expressed as a higher return to education in large CZs. We use our estimates of the basic AKM model to construct average person effects and CZ effects separately for more- and less-educated workers in each CZ, as in equation (11).

Appendix Table 5 reports the variance decompositions of our AKM estimates for the two education samples. At the individual level, they are similar to each other and to our pooled results in Table 2, though the variance of person effects is somewhat higher for the high-education sample. At the CZ level, the differences are more notable. The between-CZ standard deviation of mean earnings is almost twice as large for more-educated workers (0.178) than for less-educated workers (0.097), and the standard deviation of CZ premiums is about 50% larger for the more educated workers. In addition, the covariance between person and CZ effects – which reflects the sorting of high-skilled workers to higher-premium CZs – plays a larger role for more educated workers: The correlation between mean person effects and (education-specific) CZ wage premiums is 0.43 for lower-educated workers and 0.75 for higher-educated workers.

As is highlighted by equations (12) and (13), a key question is whether the estimated CZ effects are approximately the same for the two education groups, or substantially different. We regress $\widehat{\Psi}_{cH}$ and $\widehat{\Psi}_{cL}$ separately on our overall CZ premia $\widehat{\Psi}_{c}$. These slopes are 1.11 and 0.742, respectively, reflecting the higher standard deviations of the more-educated group's CZ premiums. The correlations, however, are extremely strong, 0.99 and 0.97.

Table 10 explores the contributions of the two main components of mean log earnings, person effects and firm effects, to the CZ size elasticity of mean log wages, separately for the two education groups. As we found in the ACS, the size gradient is much larger for more educated workers. The decomposition shows that this reflects the contributions of both person effects and CZ premiums. Differences in mean person effects by CZ size are larger for more educated workers, with a size elasticity of 0.058 versus 0.017 for less educated workers. Likewise, the mean establishment effects for more educated workers have a size elasticity of 0.043 versus 0.020 for the less educated. Columns 5 and 6 report the difference in coefficients between the two groups. This is 0.063 overall, with approximately two-thirds attributable to person effects and one-third to a higher size elasticity of the average establishment effects for more educated workers. In other words, most of rise in the college-high school wage gap in larger cities is attributable to the increased selectivity of the college-educated workforce in these cities.

Nevertheless, there is still a non-trivial role for differences in CZ premiums. We can again investigate the role that industry composition plays in this, using the decomposition of the CZ-level earnings gap between high- and low-education workers, $\bar{y}_{cH} - \bar{y}_{cL}$ from equations (12) and (13). Table 11 presents present regressions of $\bar{y}_{cH} - \bar{y}_{cL}$ and each of its components on mean log earnings in the CZ and then on the log of the CZ size. Focusing first on columns 1-2, we see that higher-wage CZs have much larger education wage gaps: A 30 log point increase in the mean wage in the CZ (approximately the difference between Houston and San Francisco) is associated with a nearly 20 log point widening of the college-high school wage gap. When we decompose the source of this divergence using equation (12), we again find that 63% is attributable to differences in person effects, leaving only about one-third to be explained by differences in the local wage premiums. When we further decompose the CZ premium component into portions reflecting industry composition and within-industry premium differences, using equation (13), we find that about two-thirds of it reflects industry composition and only one-third is attributable to differences in premiums by education within industries. Results are similar for the size gradient in columns 3-4.

It may seem counter-intuitive that industry composition has such a small effect on average CZ wage premiums across higher-wage or larger CZs (Table 9) yet explains 20-25% of the differences in average premiums for high and low education workers across higher-wage or larger CZs (Table 11).[31] The explanation is in how the sorting of higher and lower-skilled workers varies across CZs. We noted in Table 3 that the correlation of worker and firm effects within CZs is much higher in high wage places and bigger places. This increased sorting has an industry component, reflected in the measure of the composition effect (term 2) in equation 13. In essence, our results suggest that there the relative number of jobs in higher-paying industries is not much different in bigger or higher average wage CZs, but that these jobs are disproportionately filled by college-educated workers in those cities, leading to a divergence in the pay premiums earned by the two groups. As noted by Dauth et al. (2022), in a model that is additive in logarithms (i.e., super-additive), this enhanced sorting raises average labor productivity, suggesting an important benefit for larger cities.

VIII. Earnings and the cost of living

To this point we have considered decompositions of nominal earnings, unadjusted for local differences in the cost of living. But there are large and persistent differences in housing costs between places (see e.g., Moretti, 2013). Many of the coastal cities identified in Figure 2 with high causal effects on earnings have relatively high housing costs. A natural question is how the causal earnings effects that

[31] Appendix Table 6 presents a variance decomposition of the education earnings gap across CZs, analogous to Table 8. It indicates that CZs where college workers are relatively more concentrated in high-wage industries tend also to have larger skill gaps between their college and non-college workers – that is, the first term in (12) is strongly positively correlated with the 2nd term in (13).

we identify relate to differences in local costs – does moving to a larger city mean an increase in real earnings, or are the increased nominal earnings offset by higher costs?

Diamond and Moretti (2023; hereafter *DM*) use detailed expenditure and price data to construct price indexes and measures of real consumption for different income and education groups in different CZs. They find that the size elasticity of the cost of living is approximately 0.04 for higher- and lower-educated workers.[32] This elasticity is smaller than the 0.102 elasticity of college graduates' nominal earnings in Table 10. Importantly, however, we noted that after accounting for the higher skills of college-educated workers in larger CZ's, the size elasticity of our preferred place effects is only 0.043 – roughly the same as the size elasticity of the cost of living. This implies that, for college workers, higher costs fully offset the causal effect of moving to a larger city on nominal earnings, yielding approximately zero change in real earnings.[33] For non-college workers, for whom the size elasticity of estimated place effects is only 0.02, the implied elasticity of real earnings with respect to city size is negative.

To investigate these issues more directly, we use rents and housing costs information from the ACS to explore how housing costs vary with CZ size. Appendix Table 7 explores several different measures of housing costs. Across each measure, the elasticity of housing costs with respect to CZ size is 0.2 or larger. Because housing typically represents at least one-third of a household's budget, these estimates imply that nominal wages would have to exhibit a size elasticity of 0.07 or more to keep up with the cost of living (somewhat greater than the 0.04 elasticity of the cost of living estimated by DM).[34] But the size elasticities of the place effects for both college and non-college workers in Table 10

[32] DM's Figure 13 reports expenditure and real consumption size elasticities for workers in 3 education groups. The difference in these elasticities is the size elasticity of the cost of living index, which is 0.042 for workers with college education, 0.041 for workers with exactly high school education, and 0.036 for high school dropouts.

[33] DM reach a very similar conclusion, comparing the size elasticity of pre-tax income without our adjustment for unobserved skill to the size elasticity of their price index. Their income measure is quite different from ours, constructed from bank deposits. Its elasticity with respect to city size is just 0.058, about half of the 0.102 that we find for individual LEHD earnings.

[34] DM find smaller housing shares than this for high-income households. The elasticity of their full price index with respect to housing expenditures ranges from 0.24 for high-income households to 0.43 for low-income households.

are well below 0.07. Thus, using a simple 1/3 spending share on housing rule, we would conclude that housing costs consume more than 100% of the nominal earnings gain that a typical worker – even a typical college graduate – obtains from moving to a larger CZ.

An interesting question, beyond the scope of this paper, is how this can be sustained. In the standard Roback (1982) framework, causal effects of places on nominal earnings imply differences in productivity. But why would workers prefer to live in high-productivity cities if they will need to give up more than all of the earnings advantage of those cities in higher housing costs? One potential explanation consistent with Roback (1982) is the presence of consumption amenities (Albouy 2011; Albouy, Cho, and Shappo 2021). Our evidence suggests that larger CZs must offer better amenities, offsetting the reductions in real wages for workers in these more productive places.

VII. Concluding Remarks

We have used an earnings model with a combination of individual worker effects and additive establishment premiums, with a geographic structure, to address longstanding questions about the impact of place on labor market outcomes in the U.S. This class of two-way fixed effect models, originated by Abowd, Kramarz and Margolis (1999), has proven very useful in answering questions about the role of firms in wage determination (Card et al., 2018). While versions of this approach have been used to study geographic patterns in France (Combes et al., 2008), Spain (de la Roca and Puga, 2017), and Germany (Dauth et al., 2022), ours is the first to apply it to the U.S., using data from the Census Bureau's LEHD program.

We show that it is important to account for firm heterogeneity in studying geographic wage premiums, as people who move between CZs are selected from firms that occupy different positions on

---

Using 0.24 and a housing price elasticity of 0.2, the elasticity of causal earnings effects with respect to CZ size would need to exceed 0.2*0.24=0.05 to compensate for higher costs.

the local job ladders of their origin and destination places. By contrast, an AKM specification with establishment effects that are then aggregated to the CZ level provides a relatively good summary of the main patterns in the data, and can be estimated by simple OLS methods without too much concern for biases arising from either the strategic timing of moves or idiosyncratic match effects in earnings that drive mobility decisions. We find some evidence of the kind of dynamic returns to “big city” experience highlighted by de la Roca and Puga (2017) but the addition of this channel has little impact on the static returns to different CZs.

A key advantage of our specification, which allows us to separate CZ mean effects into CZ-industry terms, is that we can carefully assess the role of industry in mediating observed place effects in average earnings. Such effects can arise in two main ways. First, there can be a pure compositional effect if some CZs have a higher fraction of high-wage industries. More subtly, there can be an interaction effect if the earnings premiums for different industries vary across places and employment is locally concentrated in industries with a larger local premium. We find very small roles for either mechanism. In fact, CZ × industry pay premiums are approximately separable: Not more than 10% of the variation in these premiums is due to CZ-industry “match” effects. Moreover, we find only small interaction effects arising from the concentration of employment in sectors with a local pay premium. Wages are higher in locally agglomerated sectors, by an amount consistent with previous work (e.g., Greenstone, Hornbeck, and Moretti 2010), but industry composition does not vary enough across CZs to generate large differences in the weighted average of pure industry effects. Importantly, these conclusions are highly robust to the definition of industry. Comparing models with only 24 industries with models with close to 300, we reach nearly identical conclusions.

As in the AKM-related literature, we measure worker skills by the worker’s fixed effect in the earnings model. We find that this measure of skill varies far more widely across CZs than a more traditional measure based on observable characteristics like education, age and gender. Consistent with work in France, Spain and Germany we find that the main explanation for high wage places is the presence of high wage people. Comparing larger and smaller CZs, for example, half of the size elasticity of mean earnings is attributable to the presence of higher skilled workers. The tendency for high-wage workers to work in high-wage places magnifies overall earnings inequality in the market as a whole and represents an important feature of locational equilibrium.

We find two interesting sources of variation in the degree of assortative matching. First, college-educated workers are more highly sorted to larger and higher-wage CZs than their less educated counterparts. Nearly all of the higher "return to college" that is observed in higher-wage (or larger) places is attributable to the presence in those places of college workers with the highest unobserved skills. The differences in sorting are consistent with Diamond (2016), though her model ignores unobserved skills and treats earnings differences as causal. Second, we find that the sorting within CZs of high wage workers into high wage industries is enhanced in larger places. This confirms a similar finding by Dauth et al. (2022) in Germany, and illustrates a potential benefit of increased market size for the overall productivity of the economy.

References

Ashenfelter, Orley. "Estimating the effect of training programs on earnings." *The Review of Economics and Statistics* 60, no. 1 (1978): 47-57.

Abowd, John M., Francis Kramarz, and David N. Margolis. "High wage workers and high wage firms." *Econometrica* 67, no. 2 (1999): 251-333.

Abowd, J. M., B. E. Stephens, L. Vilhuber, F. Andersson, K. L. McKinney, M. Roemer, and S. Woodcock. "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators." In: T. Dunne, J. B. Jensen, and M. J. Roberts (Eds.), *Producer dynamics: New evidence from micro data*, vol. 68 (2009): 149-230. Chicago, IL: University of Chicago Press.

Di Addario, Sabrina L., Patrick M. Kline, Raffaele Saggio, and Mikkel Sølvsten. "It ain't where you're from, it's where you're at: Hiring origins, firm heterogeneity, and wages." *Journal of Econometrics* 233, no. 2 (2023): 340-374.

Albouy, David. "Are big cities bad places to live? Estimating quality of life across metropolitan areas." National Bureau of Economic Research working paper no. w14472, revised 2011.

Albouy, David, Heepyung Cho, and Mariya Shappo. "Immigration and the pursuit of amenities." *Journal of Regional Science* 61, no. 1 (2021): 5– 29.

Andrews, Martyn J., Len Gill, Thorsten Schank, and Richard Upward. "High wage workers and low wage firms: Negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171, no. 3 (2008): 673-697.

Autor, David H. “Work of the past, work of the future.” *American Economic Association Papers and Proceeedings* 109 (2019): 1-32.

Baum-Snow, Nathaniel, and Ronni Pavan. "Understanding the city size wage gap." *The Review of Economic Studies* 79, no. 1 (2012): 88-127.

Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. "Productive cities: Sorting, selection, and agglomeration." *Journal of Political Economy* 122, no. 3 (2014): 507-553.

Blinder, Alan S. "Wage discrimination: Reduced form and structural estimates." *Journal of Human Resources* 8, no. 4 (1973): 436-455.

Card, David, Jörg Heining, and Patrick Kline. "Workplace heterogeneity and the rise of West German wage inequality." *The Quarterly Journal of Economics* 128, no. 3 (2013): 967-1015.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics* 36, no. S1 (2018): S13-S70.

Card, David, Jesse Rothstein, and Moises Yi. “Location, location, location.” Unpublished manuscript, March 2022.

Card, David, Jesse Rothstein, and Moises Yi. “Industry wage differentials: A firm-based approach.” Unpublished manuscript, July 2023.

Chetty, Raj, and Nathaniel Hendren. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics* 133, no. 3 (2018a): 1107-1162.

Chetty, Raj, and Nathaniel Hendren. "The impacts of neighborhoods on intergenerational mobility II: County-level estimates." *The Quarterly Journal of Economics* 133, no. 3 (2018b): 1163-1228.

Ciccone, Antonio, and Robert E. Hall. "Productivity and the density of economic activity." *The American Economic Review* 86, no. 1 (1996): 54-70.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. "Spatial wage disparities: Sorting matters!" *Journal of Urban Economics* 63, no. 2 (2008): 723-742.

Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Suedekum. “Matching in cities.” *Journal of the European Economic Association* 20, no. 4 (2022): 1478–1521.

Davis, Donald R. and Jonathan I. Dingel. “A spatial knowledge economy.” *American Economic Review* 109 no. 1 (2019): 153-170.

Davis, Donald R. and Jonathan I. Dingel. “The comparative advantage of cities.” *Journal of International Economics* 123 (2020): 103291.

Diamond, Rebecca. "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000." *American Economic Review* 106, no. 3 (2016): 479-524.

Diamond, Rebecca, and Enrico Moretti. “Where is standard of living the highest? Local prices and the geography of consumption.” National Bureau of Economic Research working paper w29533, revised January 2023.

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

Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny. "Spatial sorting." *Journal of Political Economy* 122, no. 3 (2014): 554-620.

Ellison, Glenn, and Edward L. Glaeser. "Geographic concentration in US manufacturing industries: a dartboard approach." *Journal of Political Economy* 105, no. 5 (1997): 889-927.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. "Sources of geographic variation in health care: Evidence from patient migration." *The Quarterly Journal of Economics* 131, no. 4 (2016): 1681-1726.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. "Place-based drivers of mortality: Evidence from migration." *American Economic Review* 111, no. 8 (2021): 2697-2735.

Gibbons, Stephen, Henry G. Overman, and Panu Pelkonen. "Area Disparities in Britain: Understanding the Contribution of People vs. Place Through Variance Decompositions*." Oxford Bulletin of Economics and Statistics* 76, no. 5 (2014): 745-763.

Glaeser, Edward L., and Joshua D. Gottlieb. "The wealth of cities: Agglomeration economies and spatial equilibrium in the United States." *Journal of Economic Literature* 47, no. 4 (2009): 983-1028.

Glaeser, Edward L., Maré, David C. "Cities and skills." *Journal of Labor Economics* 19 no. 2 (2001): 316–342.

Gould, Eric D. "Cities, workers, and wages: A structural analysis of the urban wage premium." *The Review of Economic Studies* 74, no. 2 (2007): 477-506.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings." *Journal of Political Economy* 118, no. 3 (2010): 536-598.

Kline, Patrick, and Enrico Moretti. "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority." *The Quarterly Journal of Economics* 129, no. 1 (2014): 275-331.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. "Leave-out estimation of variance components." *Econometrica* 88, no. 5 (2020): 1859-1898.

Marshall, Alfred. *Principles of Economics*. 8th Edition. Macmillan, London (1920).

Moretti, Enrico. "Estimating the social return to higher education: Evidence from longitudinal and repeated cross-sectional data", *Journal of Econometrics* 121, no. 1-2 (2004a): 175-212.

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

Moretti, Enrico. "Local labor markets." In Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics*, vol. 4B. (2011): 1237-1313. North-Holland, Amsterdam**.**

Oaxaca, Ronald. "Male-female wage differentials in urban labor markets." *International Economic Review* 14, no. 3 (1973): 693-709.

Roca, Jorge de La, and Diego Puga. "Learning by working in big cities." *The Review of Economic Studies* 84, no. 1 (2017): 106-142.

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

Tolbert, Charles M., and Molly Sizer. “U.S. commuting zones and labor market areas: A 1990 update.” U.S. Department of Agriculture, Economic Research Service, staff paper number AGES-9614 (1996).

Vilhuber, Lars. “LEHD Infrastructure S2014 files in the FSRDC.” U.S. Census Bureau, Center for Economic Studies Discussion Paper CES 18-27 (May 2018).

Vilhuber, Lars and Kevin McKinney. “LEHD Infrastructure in the Census RDC - Overview.” U.S. Census Bureau, Center for Economic Studies Discussion Paper CES 14-26 (June 2014).

Figure 1. Mean earnings before and after a change of CZs, by change in CZ mean earnings

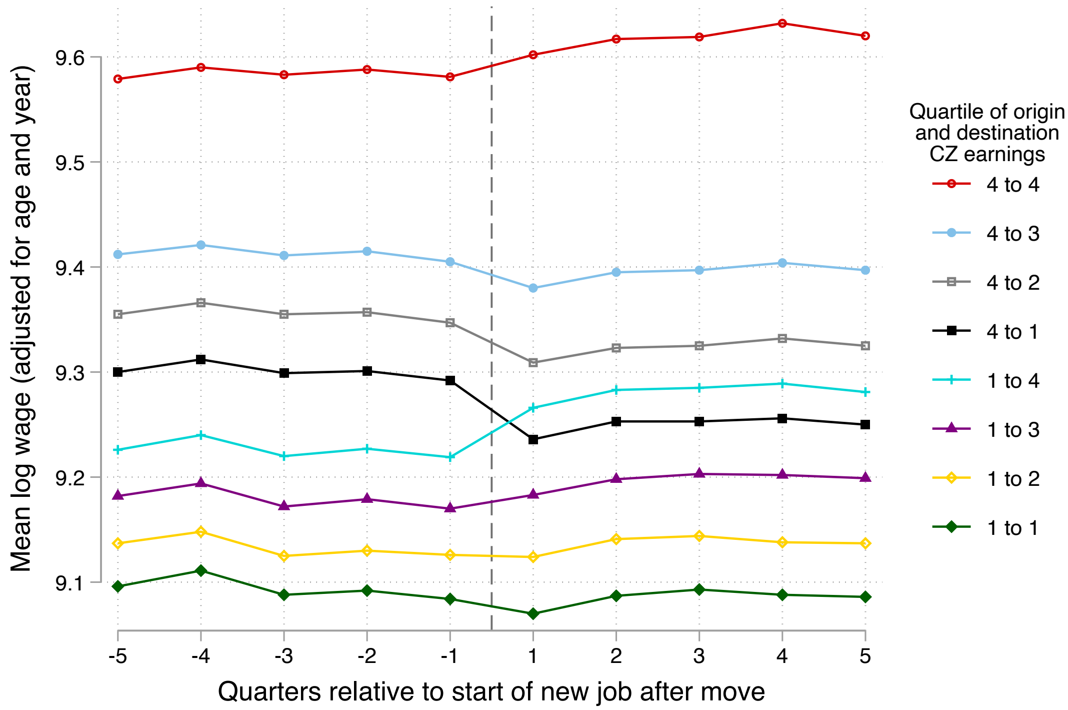

Notes: Figure shows event-time means for workers who move between CZs, separately for different quartiles of the origin and destination CZ mean earnings. Quarter -1 represents the last full quarter (i.e., the second-to-last observed quarter) in the origin CZ, and quarter +1 represents the first full quarter (second observed quarter) in the destination CZ. We allow up to four quarters of non-employment between the two observed spells. Sample is limited to workers who move between CZs only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Figure 2. Commuting zone mean wages (ACS) and CZ earnings premiums (LEHD)

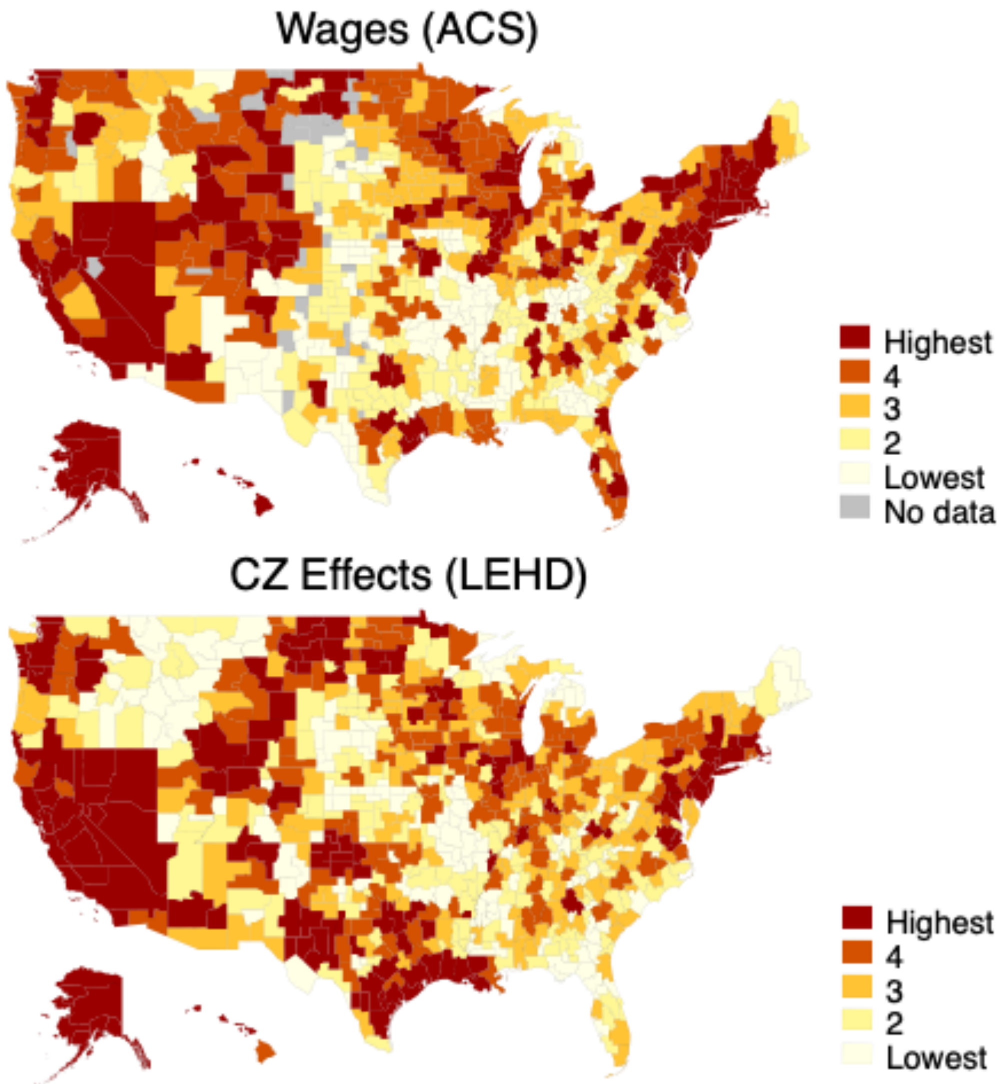

Notes: Upper map classifies CZs into quintiles based on mean hourly wages in the ACS. CZs that cannot be identified in the ACS are marked in grey. Lower map classifies into quintiles based on AKM-based CZ earnings premiums.

Figure 3. Mean earnings before and after a change of CZs, by change in CZ earnings premium

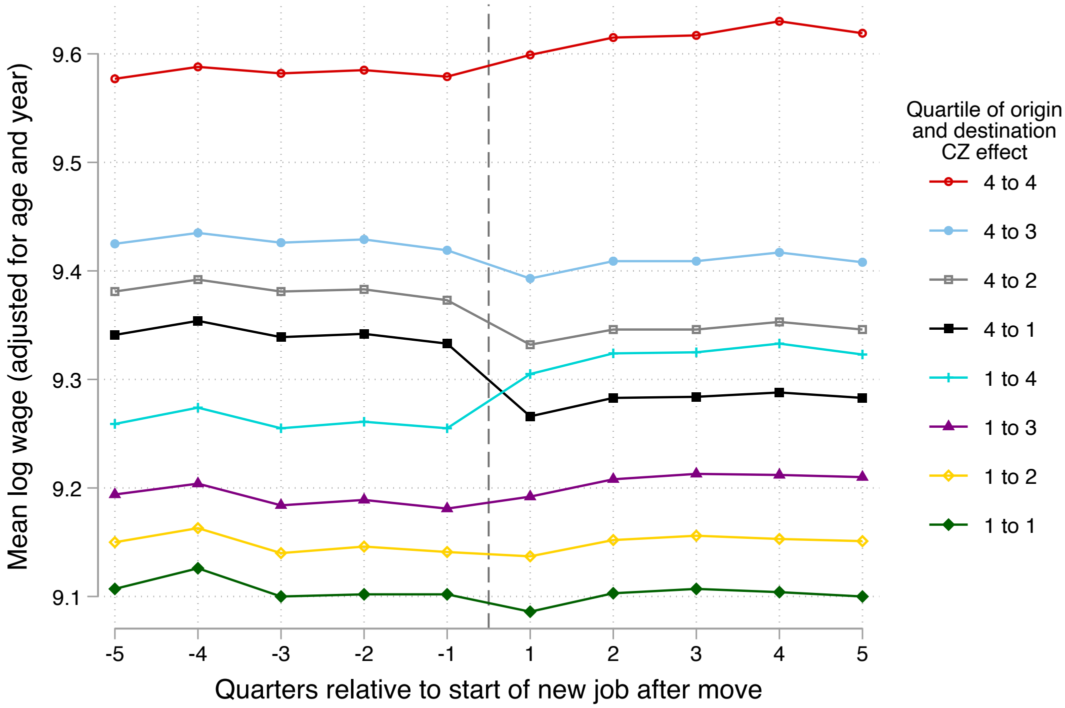

Notes: Figure shows event-time means for workers who move between CZs, separately for different quartiles of the origin and destination AKM-based CZ earnings premiums. Quarter -1 represents the last full quarter (i.e., the second-to-last observed quarter) in the origin CZ, and quarter +1 represents the first full quarter (second observed quarter) in the destination CZ. We allow up to four quarters of non-employment between the two observed spells. Sample is limited to workers who move between CZs only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Figure 4. Change in earnings components of CZ movers from last pre-move quarter to first post-move quarter, by change in CZ premium

Notes: CZs are classified into 20 vingtiles based on their AKM-based earnings premiums. We then classify CZ movers based on the change in mean premium from their origin to their destination vingtile. Figure shows the average change in the indicated earnings components from the last full quarter in the origin CZ to the first full quarter in the destination CZ for each of the resulting 400 cells. Slopes correspond to the unweighted best linear fit line fit to these 400 points. Sample is limited to workers who move between CZs only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Table 1: Characteristics of Samples Derived from Longitudinal Employer-Household Dynamics (LEHD) Data Base

	Estimation Sample	Subgroups of Estimation Sample:		Event study samples:	
		CZ stayers	CZ movers	Firm movers	CZ movers
	(1)	(2)	(3)	(4)	(5)
ln(quarterly earnings)	9.43	9.43	9.44	9.51	9.55
(standard deviation)	(0.60)	(0.59)	(0.60)	(0.63)	(0.65)
Mean Age	42.4	43.6	39.5	40.9	40.1
Fraction Female	0.47	0.49	0.43	0.46	0.43
Fraction Foreign Born	0.16	0.17	0.14	0.15	0.13
Fraction Hispanic	0.13	0.14	0.12	0.12	0.10
Number of CZ's during sample period:					
1 CZ	0.73	1	0	0.67	0
2 CZ's	0.20	0	0.75	0.33	1
3+ CZ's	0.07	0	0.25	0.00	0
Number of 4-digit industries during sample period:					
1 Industry	0.65	0.68	0.57	0.62	0.99
2 Industries	0.25	0.23	0.30	0.37	0.01
3+ Industries	0.11	0.09	0.13	0.01	0.00
Number of Quarters Observed	25.9	26.8	23.3		
Number P-Q obs (millions)	2,523	1,835	687	193	67
Number Persons (millions)	112	79	34	19	7

Notes: Sample includes person-quarter (PQ) observations for indivduals age 22-62 with at least 8 quarters of employment in the LEHD 2010Q1 to 2018Q2. Quarterly observations for individuals with multiple employers are excluded, as are the first and last (transitional) quarters of any spell with the same employer, quarters for which industry or location information is missing, and quarters with earnings less than $3,800. Event study samples use 5 quarters of earnings before a move and 5 quarters after, and are limited to people with a single move (between firms in column 4, or between CZs with no change in industry in column 5) in that window.

Table 2: Summary of Estimated AKM Model

	Person-quarter level		CZ level	
	Std. Dev. or Correlation	Variance Share	Std. Dev. or Correlation	Variance Share
	(1)	(2)	(3)	(4)
Log earnings or mean log earnings	0.653	1.000	0.147	1.000
Variance components (std. deviations in odd-number columns, variance shares in even-numbered columns)				
Person effects	0.524	0.644	0.081	0.303
Firm effects	0.223	0.117	0.079	0.293
Covariate index (Xβ)	0.135	0.043	0.006	0.002
Residual	0.232	0.127	0.000	0.000
Covariance components (correlations in odd-number columns, variance shares in even-numbered columns)				
Person & firm	0.182	0.100	0.642	0.382
Person & covariate index	-0.100	-0.033	0.228	0.010
Firm & covariate index	0.020	0.003	0.245	0.010

Notes: Table shows variance decompositions based on equation (4). Columns 1-2 pertain to the variance of individual quarterly earnings. Columns 3-4 pertain to the variance of mean earnings by CZ. Entries in columns 1 & 3 for "variance components" are standard deviations of the earnings components indicated in row headings; for "covariance components" they are the estimated correlations of the indicated components. Entries in columns 2 & 4 are variance shares explained by the variance or covariance components.

Table 3. CZ characteristics and AKM estimates

	Explanatory variable (CZ level)		
	Mean log earnings	Log size	Fraction college
	(1)	(2)	(3)
Mean log earnings	1	0.075	1.683
		(0.008)	(0.149)
Basic decomposition			
Mean person effect (skill composition effect)	0.499	0.040	0.982
	(0.021)	(0.007)	(0.069)
Mean establishment effect (CZ place effect)	0.489	0.034	0.664
	(0.019)	(0.003)	(0.099)
Share of gradient due to skill composition	49%	45%	39%
Measures of dispersion in skill composition			
SD of person effects	0.240	0.025	0.466
	(0.022)	(0.002)	(0.047)
Share of workers in bottom decile	-0.087	-0.005	-0.202
	(0.012)	(0.003)	(0.021)
Share of workers in top decile	0.213	0.019	0.414
	(0.009)	(0.002)	(0.034)
Within-CZ skill-match			
Correlation of person & firm effect within CZ	0.451	0.039	0.850
	(0.025)	(0.004)	(0.057)

Notes: Table shows regressions of mean log earnings and other CZ-level measures on CZ mean log earnings, log size, or the fraction of workers with some college or more. Regressions are weighted by CZ size, measured as the number of person-quarter observations in our sample in the CZ.

Table 4: Mean Log Earnings and Components of Earnings for Groups of CZ's

	Ranked By CZ Effect (Ψ_c)					Grouped By Quintile of CZ Effect (Ψ_c; 1=lowest)				
		Top 10			Bottom 10					
	All (1)	Large Urban (2)	Resource Intensive (3)	Middle Range (4)	(5)	1 (6)	2 (7)	3 (8)	4 (9)	5 (10)
Log Earnings	9.26	9.72	9.54	9.29	9.13	9.14	9.20	9.25	9.31	9.42
	(0.12)	(0.08)	(0.05)	(0.10)	(0.07)	(0.06)	(0.05)	(0.06)	(0.07)	(0.13)
Person Effect	9.02	9.24	9.07	9.04	9.00	8.99	9.00	9.02	9.04	9.05
	(0.07)	(0.03)	(0.03)	(0.07)	(0.06)	(0.05)	(0.05)	(0.06)	(0.07)	(0.09)
CZ effect	0.18	0.41	0.41	0.18	0.07	0.09	0.14	0.17	0.21	0.31
	(0.09)	(0.06)	(0.03)	(0.05)	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)	(0.09)
Covariates	0.06	0.07	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
	(0.01)	(0.00)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)

Notes: Analyses are weighted by CZ size. Mean and standard deviations (in parentheses) are based on the across-CZ distribution. Columns 2-5 include only CZs with at least 25,000 people.

Table 5: Dynamic specifications, estimated on younger workers

	AKM models estimated on young subsample		
	No dynamics	Dynamic models, large CZs defined: 10 Largest CZ's	Dynamic models, large CZs defined: 25 Largest CZ's
	(1)	(2)	(3)
A. Regression of CZ effect from alternative specification on CZ effect from baseline specificati			
Std. dev. of estimated CZ effects	0.078	0.077	0.078
Slope coefficient from regression on	0.96	0.94	0.95
CZ effects from main model	(0.03)	(0.03)	(0.03)
R-squared coefficient	0.94	0.94	0.94
B. Estimated effects of work experience from dynamic specification:			
Years of work experience in any CZ		0.017	0.014
		(0.001)	(0.001)
Years of work experience in large CZ		0.016	0.016
		(0.007)	(0.004)
Years of work experience in any CZ		0.007	0.002
× in large CZ		(0.007)	(0.004)
Years of work experience in large CZ		-0.008	-0.002
× in large CZ		(0.001)	(0.001)

Notes: Panel A shows regressions of estimated CZ effects from alternative models, all estimated using only individuals under age 26 in 2010, on the CZ effects from our main specification. Number of observations is 741 (regressions weighted by CZ size). Robust standard errors in parentheses. Panel B shows estimated regression coefficients for experience terms from the dynamic specifications. In column 2, "large CZs" are the 10 largest CZ's, and "large CZ experience" is previous quarters of work in these 10 CZs. In column 3, "large CZs" are the 25 largest CZ's, with "large CZ experience" defined commensurately. Standard errors, clustered by the first CZ in which a person is observed in our sample, in parentheses.

Table 6. Comparisons of Estimated CZ Pay Premiums from Alternative Specifications to Estimates from Preferred Approach

	Preferred model	Models with Person and Aggregated Place Effects		Cross-sectional models fit to ACS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Alternative model controls for:							
Year effects and age/experience	X	X	X		X	X	X
Educ., experience, gender, race					X	X	X
Region of origin & field of study						X	X
Detailed industry							X
Person fixed effects	X	X	X				
Firm fixed effects	X						
CZ fixed effects		X					
Industry-by-CZ fixed effects			X				
Standard deviation of CZ effects	0.079	0.064	0.065	0.141	0.117	0.111	0.106
Regression of CZ effects from alt. model on preferred estimates	1.00	0.77	0.79	1.43	1.33	1.27	1.18
		(0.02)	(0.02)	(0.11)	(0.05)	(0.06)	(0.06)
R^2 (adjusted)	1.000	0.919	0.945	0.656	0.816	0.813	0.777

Note: Preferred model construct CZ effects as averages of firm effects from AKM specification, as in equation (3). Regressions are of CZ effects from alternative models on CZ effects from preferred model, and are weighted by the number of person-quarter observations in the CZ (in columns 2-3) or by the weighted number of ACS wage observations in the CZ (in columns 4-7). Alternative models in columns 2-3 are fit to LEHD with person effects and CZ or CZ-by-industry effects, and include the same year and age effects as the preferred model. Alternative models in columns 4-7 are fit to the ACS with observed control variables and CZ effects. Robust standard errors in parentheses.

Table 7: Simple Linear Regression Models for Estimated CZ-Industry Effects

	CZ Effects only	Industry Effects only	Models with CZ and Industry Effects: No other controls	Plus local emp. share of industry	Plus log of local emp. share of industry
	(1)	(2)	(3)	(4)	(5)
Share of CZ emp. in industry	--	--	--	0.451	--
				(0.033)	
Log share of CZ emp. in industry	--	--	--	--	0.044
					(0.002)
R-squared	0.326	0.570	0.873	0.885	0.892
Root Mean Squared Error	0.117	0.091	0.051	0.048	0.047

Notes: Table shows goodness of fit and estimated regression coefficients from regression of estimated CZ-by-2-digit-industry effects on CZ indicators (column 1), industry indicators (column 2), and CZ and industry indicators (columns 3-5). Models in columns 4-5 include the share (or log share) of CZ employment in the industry. Models are fit to person-quarter data after assigning estimated CZ-industry effect and all control variables to each person-quarter observation.

Table 8: Decomposition of Variance of Average CZ Earnings Premium

	2-digit industries (2)	4-digit industries (3)
Standard Dev. of Average CZ premium	0.079	0.079
Decomposition (variance shares):		
Var(Average Earnings Premium)	0.942	0.913
Var(Composition Effect)	0.025	0.026
Var(Interaction Effect)	0.022	0.044
Cov(Earnings Premium, Composition Effect)	0.015	0.064
Cov(Earnings Premium, Interaction Effect)	-0.006	-0.054
Cov(Composition Effect, Interaction)	0.001	0.008

Notes: Table shows decomposition of the variance of estimated average CZ wage premium, based on equation (9) in text. Decompositions uses 24 2-digit industries (column 1) or 312 4-digit industries (column 2) to define CZ-by-industry effects.

Table 9. Decomposition of Effect of CZ Characteristics on CZ Earnings Effects

	CZ Characteristic (explanatory variable)		
	Mean log earnings	Log size	College share
	(1)	(2)	(3)
Mean log earnings	1	0.075	1.683
		(0.008)	(0.149)
Basic decomposition			
Mean person effect	0.499	0.040	0.982
	(0.021)	(0.007)	(0.069)
CZ earnings effect	0.489	0.034	0.664
	(0.019)	(0.003)	(0.099)
Decomposition of CZ earnings effect (based on equation 9)			
Average earnings premium	0.475	0.037	0.733
(term 1 in equation 9)	(0.019)	(0.003)	(0.089)
Industry composition effect	0.018	0.0000	-0.010
(term 2 in equation 9)	(0.006)	(0.0004)	(0.013)
Interaction effect	-0.004	-0.003	-0.059
(term 3 in equation 9)	(0.008)	(0.001)	(0.015)

Notes: Table shows regressions of mean log earnings and components of the decomposition given by equations (6) and (9) on CZ mean log earnings, log size, or the fraction of workers with some college or more. Regressions are weighted by CZ size.

Table 10. Decomposition of CZ Size Elasticity of Mean Wage by Education Group

	Less educated workers		More educated workers		Difference, more educated - less educated	
	Coefficient	Share	Coefficient	Share	Coefficient	Share
	(1)	(2)	(3)	(4)	(5)	(6)
Log earnings	0.039	100%	0.102	100%	0.063	100%
	(0.006)		(0.008)		(0.010)	
Components						
Mean person effect	0.017	43%	0.058	57%	0.041	65%
(skill composition effect)	(0.005)		(0.006)		(0.008)	
Mean establishment effect	0.020	51%	0.043	42%	0.023	36%
(CZ place effect)	(0.002)		(0.003)		(0.004)	

Notes: Table shows regressions of mean log earnings by education group (high school or less vs. some college or more) and components of the decomposition given by equation (6) on CZ mean log earnings, log size, or the fraction of workers with some college or more. Regressions are weighted by the number of person-quarter observations in the relevant education group. Column 5 represents the difference between column 3 and column 2; standard errors assume estimates are independent across the education groups.

Table 11. Decomposition of Variation in the CZ-level Return to Education

	Explanatory variable:			
	Mean log earnings in CZ		Log CZ size	
	Coefficient	Share	Coefficient	Share
	(1)	(2)	(3)	(4)
Overall earnings gap between more and less educated workers	0.628 (0.049)	100%	0.060 (0.003)	100%
Components				
Person effects	0.398 (0.034)	63%	0.041 (0.002)	68%
Establishment effects (CZ wage effect)	0.233 (0.017)	37%	0.020 (0.001)	33%
Covariates	-0.003 (0.003)	-1%	-0.001 (0.000)	-1%
CZ wage effect components				
Relative wage premium	0.074 (0.009)	12%	0.007 (0.000)	11%
Composition	0.165 (0.011)	26%	0.013 (0.001)	22%
Interaction	-0.006 (0.002)	-1%	0.000 (0.000)	0%
Sum of relative wage premium and interaction	0.068 (0.007)	11%	0.006 (0.000)	11%

Notes: Table shows regressions of the difference in mean log earnings of college and non-college workers in a CZ, or components of the decomposition given by equations (12) and (13), on CZ mean log earnings and log size. Regressions are weighted by CZ size.

Appendices to "Location, location, location"
David Card, Jesse Rothstein, and Moises Yi
August 2023

Appendix A. Samples

1. ACS Worker Level Files

Our ACS sample is formed by pooling public-use microdata for 2010-2018. We focus on individuals aged 18-62 with at least one year of potential experience, constructing an estimated hourly wage from data on total earnings in the previous year, hours per week, and information on weeks worked. To limit the effect of outliers we Winsorize the hourly wage at $5 and $500.

We assign local labor markets based on the 1990 CZ boundaries developed by Tolbert and Sizer (1996), which define 741 CZs, each comprised of one or more whole counties. The lowest level of geography in the ACS is the public use micro area (PUMA), which can contain multiple counties (or parts of counties) in sparser areas. We use the fractions of people in each PUMA who lived in each county in the 2000 Census (for the 2010 and 2011 ACS) or the 2010 Census (for the 2012+ ACS) to probabilistically allocate respondents to counties and CZs. We pool all CZ's in Alaska, yielding a total of 691 commuting zones. We note that the size distribution of CZs is highly skewed: The 50 largest CZs have nearly 60% of workers, and the 200 largest have over 85%.

Our sample has 11.7 million workers, providing relatively large samples for even modest-sized CZ's. We have around 10,000 observations for the CZ's ranked at roughly 200th in size (e.g., Binghamton, NY; Morgantown, WV; and Byron, TX), and 3,000-3,500 for CZ's ranked at roughly 400th in size.

The mean of log nominal hourly wages in our ACS sample is 2.863 (about $17.50), and the (weighed) standard deviation across CZ's is 0.141. The $\pm 2\sigma$ "Lester range" is therefore 56 log points or about 75%. The coefficient of a (weighted) regression of mean log hourly wages on log size is 0.068 with a robust standard error of 0.01 (Appendix Table 1). Log size explains about 50% of the cross-CZ variance in mean log wages. It is also highly correlated with the share of college-educated workers, the fraction of immigrants, and the share of white non-Hispanic workers.

We construct three measures of skill-adjusted CZ effects using a (weighted) regression of log wages on CZ effects and controls. In the first (basic) model, the model includes 18 individual controls (education, a quartic of experience interacted with gender, and dummies for race/ethnicity, interacted with gender, along with year effects. This model has an R-squared of 0.3183, and a root mean squared error of 0.6029.

Our second model generalizes this by allowing dummies for individual years of completed education (interacted with gender, immigrant status, and whether an immigrant's years in the US is less than their potential experience + 5 (implying that the probably completed some years of schooling in the U.S.), as well as dummies for 3 major immigrant source regions (Latin America, Asia, and Europe/Canada/Australia/New Zealand) and controls for years in the U.S. for immigrants (interacted with source region). This model has an R-squared of 0.3490, and a root mean squared error of 0.5892.

Our third model generalizes the second model by adding controls for field of degree for people with a BA or higher education (in 16 main categories), interacted with gender, and controls for detailed industry, based on Census industry codes used in the ACS (a total of 267 codes). We adjust the coding of Census industry across years of the ACS to the 2018 coding system. This model has an R-squared of 0.400, and a root mean squared error of 0.5654.

We also fit two extra versions of this third model, using data on people with exactly 12 and exactly 16 years of education. The high school sample has 3.755 million observations; the R-squared of the model is 0.253 and the root mean squared error is 0.5296. The college sample has 2.516 million observations; the R-squared of the model is 0.302 and the root mean squared error is 0.5875.

Our wage models are limited to people who were employed in the reference year. As a final exercise, we fit versions of the second model to the full ACS sample, including non-workers, where the dependent variables are an indicator for annual employment and the continuous number of hours worked in the year (including non-workers as zeros). We use these below to explore geographic variation in employment and hours.

2. ACS Household Level File (For House Price/Rent Analysis)

Our ACS household file is formed from the household records of the 2014-2018 5-year ACS public use sample. We extract all non-group-quarters records, providing a sample of 6.804 million household records. We assign CZs using the same procedures described above, using PUMA codes and information from the 2010 Census to probabilistically allocate households to counties and CZs. The sample has 692 separately identified CZs. We adjust reported property values and rents to real values using the deflator variable ("ADJINC") provided in the ACS. We set reported property values in the interval of $1,000 to $4,500 to $4,500, and drop observations with reported value below $1,000. Similarly we drop rental values of less than $100 per month. This leaves 4.374 million observations on property values (for households that own their home), and 1.836 million observations on rents (for households that rent). The mean of log housing value is 12.14 (implying a geometric mean of $187,213) with a standard deviation of 1.03. The mean of log monthly rents is 6.70 (implying a geometric mean of $815) with a standard deviation of 0.67.

To obtain adjusted property values and rents, we fit simple models for the logs of property values and log of monthly rents that include CZ effects and housing unit characteristics. For property values, the controls include: dummies for whether the unit is a mobile home, a single attached home, or a unit in buildings in 6 different size ranges, controls for the number of bedrooms (5 dummies), the log of the total number of rooms, dummies for the year the unit was built (in 22 ranges), and a dummy for whether the homeowner has a mortgage. This model has an R-squared of 0.50 and a root mean squared error of 0.70. For rents the controls include the same characteristics, plus indicators for whether electricity, gas, water, other fuels, or meals are included in the rent (a total of 5 indicators). This model has an R-squared of 0.34 and a root mean squared error of 0.52.

3. LEHD

We discuss here a few details of the LEHD sample that were not discussed fully in the text.

One is that the unemployment insurance records from which the LEHD is constructed identify the firm at which a worker is employed and the state, but not the specific establishment. Only one state reports establishment locations. The Census Bureau uses this state's data to train a model that it then uses to impute establishment assignments probabilistically to data from other states, using the distance between the worker's residential location and each of the establishments of the employing firm (Vilhuber 2018). Ten imputations are reported for each PEQ. When the firm has only a single establishment within reasonable commuting distance of the worker, all ten of these are identical.

Our analysis uses establishments both for the AKM model and to assign CZs and industries. We classify PEQs based on whether there is variation in the implied CZ or industry across the ten imputations. We present an analysis below of the no-uncertainty sample.

Section VII segments the sample by worker education (high school or less, vs. some college or more). Although education is imputed for most workers in the LEHD, these imputations carry some error. As an alternative, we take advantage of linkages that the Census Bureau has created between LEHD records and the 2001-17 ACS files. Any individual captured in one of these survey samples was asked about his or her completed education. Our education analyses restrict attention to observations with linked education information from one of these surveys, using only survey responses when respondents were at least 30 years old. This yields about 15% of the full LEHD sample with education information.

Appendix B. Additional results

1. Additional ACS results

Section II of the paper discusses four stylized facts that we derive from our analysis of the ACS sample introduced in Appendix A: (1) mean wages vary widely across CZs; (2) only a modest share of this variation is explained by differences in observed characteristics of the workers in different CZs; (3) mean wages are higher in larger CZs; and (4) pay premiums for working in larger cities or in cities with higher average wages are higher for better-educated workers. We present here evidence in support of these claims.

Appendix Figure 2 compares mean log wages in each CZ, on the x-axis, to estimated CZ wage effects from a model that controls for worker observables, on the y-axis. This is our third, most detailed observational model discussed above. The wide spread of points on the x-axis, with a standard deviation of 0.14 and an interquartile range of 0.2, establishes the first stylized fact. The slope of the fitted line on the figure, 0.69, demonstrates that less than one-third of this variation is explained by observed worker characteristics. The standard deviation of the adjusted CZ effects is 0.106, about 75% as large as the variation in unadjusted mean wages. In Appendix Figure 2 we plot the adjusted CZ effects against the mean wages in each CZ. The slope is 0.692, suggesting that only about 30% of the variance of CZ wages can be explained by even a rich set of observed worker and job characteristics.

Appendix Figure 3 compares unadjusted CZ mean wages and adjusted CZ effects to log CZ size. Both unadjusted means and adjusted CZ effects are higher in large CZs, with elasticities of 0.068 and 0.056, respectively. Appendix Table 1 presents regressions of unadjusted and adjusted wages on CZ size, as well as similar estimates for annual earnings. Elasticities of earnings with respect to size are very similar to those of wages, suggesting that labor supply does not vary dramatically across CZs (at least among those who work).

Appendix Figure 4 compares our adjusted CZ effects fit separately for workers with 12 and with 16 years of education to the estimates from the model that pools all education groups. (All models are normalized to mean zero in the average CZ.) The slope of the 12-year estimates with respect to the pooled estimates is 0.906, while that for the 16-year estimates is 1.209. The implication is that the return to education is larger in higher-premium CZs.

Finally, we used information on annual employment and hours in the ACS to study geographic variation in labor supply. Appendix Table 2 presents results. Rows 1 and 3 show models relating CZ mean labor supply to CZ size (column 1), the unadjusted mean log wage in the CZ (column 2), or the adjusted log wage from the second ACS specification described above (column 3). Both labor supply measures are positively related to CZ size and CZ unadjusted and adjusted wages. In rows 2 and 4, we replace the unadjusted means of the dependent variables with adjusted CZ effects, using the same specification. The adjusted labor supply measures remain positively related to unadjusted mean wages, though the relationship is much weaker. They are not positively related either to CZ size or to adjusted CZ mean wages. Columns 4 through 6 repeat the exercise separately by gender, with similar results for men and women.

We conclude that there is no evidence that employment varies systematically at either the intensive or extensive margin with CZ pay premiums.

2. Evaluating the AKM specification

In this section, we present a number of analyses aimed at validating the restrictions of the AKM specification (2) – exogenous mobility with respect to the error term, and additive separability of the person and establishment effects. These analyses replicate and extend analyses proposed by Card et al. (2013), and are intended in part to allow for comparisons between U.S. data and results seen in other countries.

We begin with an analysis of additive separability. We examine the residuals from equation (2), looking for evidence that the mean residuals for high or low skilled workers (classified by the estimated value of α_i) are larger in magnitude for jobs in high or low premium establishments (classified by the estimated values of δ_f). Appendix Figure 5 shows mean residuals by decile of α_i crossed with decile of δ_f. Systematic patterns might point to violations of the assumption that log wages are additively separable in person and firm effects. For example, if log wages were determined by a function $g(\alpha_i, \delta_f)$ with a positive cross-partial derivative, we would expect residuals to be positive at the far right corner of the Figure and negative at the near left corner. We do not see this. Generally, mean residuals are quite small. The main pattern we see is a positive average residual at the near left corner, when the lowest-skill workers are seen at the lowest-premium establishments, and corresponding negative residuals in the near right and far left corners, very high (low) skill workers at very low (high) premium establishments. (Note that residuals must average to zero within each row and column.) This could be consistent with a $g()$ function with a negative cross-partial derivative, or, perhaps, with minimum wage effects that put a floor on the very lowest wages. In any event, the deviation is quite small even at the extremes, and is almost entirely concentrated in the lowest decile of person and firm effects. Thus, violations of additive separability appear to be small.

The main text presented event study plots of mean earnings in the quarters leading up to and following moves across commuting zones (e.g., Figure 3). Evidence that workers who moved had systematically different trends or shocks prior to the move would tend to contradict the exogenous mobility assumption. Appendix Figure 6 presents a different view of this event study, this time plotting the evolution of the equation (2) *residuals*, rather than total earnings, around the move. We see some evidence here that residuals dip by as much as 2% in the first quarter after a move to a new CZ, with the larger dips among the workers moving up the most in the CZ premium distribution. However, this tip is quickly recovered.

Appendix Figure 7 presents a similar event study, this time centered around moves across *establishments*. As before, we show series for moves out of the top- and bottom-quartile of establishments classified by δ_f. Patterns are generally similar to those for CZ moves, though more extreme reflecting the greater variance of establishment effects than CZ means. Earnings are generally quite flat prior to moves, with perhaps a bit of indication that 4-to-1 downward movers had already been experiencing very slow earnings declines prior to the move. Levels are consistent with expectations: Workers in quartile-4 establishments have much higher earnings than those in quartile-1 establishments, while within the quartile-4 group those who will eventually move to lower-premium establishments already have lower earnings prior to the move (which our model interprets as indicating lower α_is) than those who will not. Earnings changes following moves are also consistent with expectations, with larger changes for those with the biggest shifts in establishment premia. As in earlier analyses, we see some indication that quarter-1 earnings are a bit depressed, with most groups rising a bit after that but little indication of differential patterns in this rise.

The right panel of Appendix Figure 7 zooms in by plotting just the residual component of earnings. With the much-enlarged scale, we can see better that the model assumptions do not seem to hold exactly. Workers who make big upward moves to higher-δ_f firms see non-trivial declines in their residuals in the first quarter following the move, suggesting that they do not receive the full expected wage increase immediately. These residuals disappear fairly quickly afterward; by the 4^{th} quarter after the move, residuals for all of the groups have settled into the same 0.05 range that they were in two quarters before the move. The ranking within this range has changed slightly, but the range is small enough not to leave much scope for match effects.

(Recall that the standard deviation of estimated establishment effects is 0.22, so a residual shift of, say, 0.03 is not large in comparison.)

Appendix Figure 8 repeats the exercise from Figure 4, this time for between-establishment rather than between-CZ moves. We divide establishment into vingtiles based on their estimated earnings premiums, and plot the average change in log earnings (removing the contribution of $X_{it}\hat{\beta}$ to obtain age-adjusted earnings) from period -1 to period 1 for movers in each of the 400 possible origin and destination cells against the predicted change in δ_f given the origin and destination vingtiles. This is a more direct test of the AKM specification than was Figure 4, as equation (2) specifies that age-adjusted earnings are $y_{it} - X_{it}\hat{\beta} = \delta_{f(i,t)} + \varepsilon_{it}$, so only systematic changes in the residual could produce deviations from the 45-degree line.

The scatterplot shows that changes in earnings are generally quite close to that line, with a slope of about 0.87. What deviations there are indicate that movers see somewhat smaller changes in wages (in either direction) than would be predicted from the change in firm effects. Panel B shows the same plot, this time with residuals. It is easier to see the deviations here: People who move to firms where the premium increases (decreases) by about 0.4 see their residuals decrease (increase) by about 0.1 in the first quarter after the move, with smaller changes for more extreme moves.

Appendix Figure 6 indicated that residual changes in the first quarter after a move are particularly extreme. When we repeat the analysis in Appendix Figure 8, panel B, using the change in residuals between the last pre-move quarter and the *fourth* post-move quarter, the slope diminishes, to just -0.08. Thus, while there is some evidence that match effects or residuals offset some of the differences in pure firm effects, the magnitude is quite small.[1]

A close examination of Appendix Figure 8 suggests that there may be some asymmetry in wage changes for movers, with a somewhat larger gap between predictions and actual wage changes for those who move to firms with higher premiums than for those who move to firms with lower premiums. We investigate this further in Appendix Figure 9. Here, we use the same 20-by-20 grid of vingtile-to-vingtile moves. For each pair (i,j) where j is a higher vingtile than i,

[1] We have conducted this analysis separately for workers in each tercile of the distribution of estimated α_i, to gauge violations of additive separability. We find quite similar slopes across terciles.

we compute the average observed earnings change from period -1 to +1 for i-to-j movers (i.e., for those who move to higher-premium firms) and for the corresponding group of j-to-i movers. Appendix Figure 9 arrays the averages for downward moves on the y-axis against the corresponding averages for upper moves on the x-axis.

The AKM specification (2) implies that these should be equal and opposite, at least approximately. (Because there is variation in δ_f within vingtiles, it is possible that the predicted changes for observed movers in the two directions are not exactly opposite.) This is not quite what we see. The average downward mover loses a bit more than the average upward mover gains. This is an indication that the AKM model does not fit the data exactly. However, the deviations remain small. Moreover, we have already seen that the first post-move quarter is somewhat anomalous. When we repeat the exercise using data from the 4th post-quarter move, the deviation is much smaller (yielding a slope of -1.11, as compared with the -1.30 in Appendix Figure 9 for the first post-quarter move).

We also show a similar symmetry test for between-CZ moves, in Appendix Figure 10. We have much less variation to work with here, as the variation in firm-level premiums is so much larger than that in CZ-level premiums. However, there is no indication here of the lack of symmetry we saw at the firm level. One interpretation is that within-CZ, between-firm moves, particularly to higher-premium firms, are more likely to depend on changes in match effects or non-wage amenities that partially offset the change in firm premiums than are between-CZ moves.

Overall, our investigation indicates that equation (2), with its implicit assumptions of additive separability and exogenous mobility, fits the data reasonably well. We do find some evidence of small violations of the exogenous mobility assumption, particularly in the first few quarters after a move, that could be incorporated in future work. However, so far as we are able to assess, these violations seem to be most severe for within-CZ moves, and wage changes surrounding between-CZ moves seem to come much closer to the AKM model ideal.

3. Public estimates of CZ wage premiums

There would be value in distributing our estimates of CZ wage premiums for use in other analyses. Unfortunately, we have not been able to disclose CZ-level statistics from the LEHD, beyond the classification into quintiles shown in Figure 2. As an alternative, we have produced approximations to the LEHD-based premiums based on public information from the ACS and disclosed regression coefficients (presented in Tables 6 and 8). These are predicted values from regressions of the LEHD premiums on three ACS variables: The estimated CZ wage premium from our first cross-sectional model, the log of the CZ size, and the share of the workforce with some college or more. The fitted values from this regression have standard deviation 0.072, so are correlated 0.90 with the LEHD estimates (and thus account for 83% of the variance of those estimates). They are unbiased predictors, so substitution of these ACS-based measures for the undisclosed LEHD measures as independent variables in regressions should not bias estimates (Hyslop and Imbens 2001).

The ACS-based CZ wage premium estimates are reported for the 50 largest CZs in Appendix Table 3, and for all CZs in a data file that will be distributed with this appendix.

4. Extended analyses

We present here a few additional results.

First, Appendix Table 4 repeats the variance decomposition of our AKM estimates from Table 2, both for the full sample shown earlier (columns 1-2) and for two subsamples. (In each case, we use an AKM model fit to the entire sample, and restrict to subsamples only in summarizing that model's results.) In columns 3-4, we show an analysis of non-movers, people observed only in a single CZ in our sample. Because the AKM model is identified from movers, with results then extrapolated to stayers, evidence that stayers differ substantially from movers would reduce confidence in the analysis. We see that the decomposition for stayers is similar to that in the full sample. In columns 5-6, we limit to person-quarter observations for which the worker is at a single-establishment firm or for which the establishments are sufficiently geographically dispersed (or sufficiently similar) that all ten imputations of the establishment assignment yield the same industry and CZ assignment. Again, this makes little difference to the results.

Next, we presented analyses in the main text of differences across education groups. We present additional results for this analysis in Appendix Table 5. As in Appendix Table 4, we fit the main AKM model with the full sample, then analyze separately the low-education and high-education subsets of this sample. (This is necessary because we observe education only for a small share of observations; limiting the AKM estimation only to those with observed education at a particular level would eliminate many of the connections needed to identify the AKM model.) We show the variance decomposition separately for the two groups. Firm effects account for a larger share of the variance of earnings among low-education workers, while among high-education workers person effects are a bit more important in variance terms.

Implicit in this analysis' reliance on a pooled AKM model is an assumption that firm premiums δ_f are the same for high- and low-education workers. To assess this, we computed worker-firm match effects as the average residual for each worker-firm combination. Insofar as firm premiums differ for the two types of workers, that difference would be captured in the match effects (e.g., with positive match effects for H workers and negative match effects for L workers). We then regressed these match effects on firm-by-education indicators. The adjusted R-squared of this regression was negative, offering no indication of differences in firm premiums by education.

Appendix Table 6 addresses uses the five-part decomposition arising from equations 12 and 13 to explore the overall variance of the CZ-level return to education. As in Table 11, we see a large role for person effects, with differences in the skill gap between college and high school workers accounting for about one-half of the variance in returns. Most of the rest is explained by two important covariance terms: 14% is attributable to covariance between person effects and the pure CZ wage premium (the 1st term in equation (13)), and 28% is attributable to covariance between person effects and industry composition. The latter effect confirms the conclusion from Table 11 that differences in the degree of sorting of higher-education workers to high-wage industries is an important source of variation in the wage gap between college and high school workers across different CZ's.

Finally, we discuss our analysis of housing price differences across CZs, reported in Appendix Table 7. We construct four CZ-level housing measures - mean log home values for homeowners and mean log rents for renters, each unadjusted and then adjusted for housing characteristics[2]) on log CZ size. Characteristics are type of unit (with five classes of apartment building size), number of bedrooms and of total rooms, year of construction, indicators for utilities included in rents (for rents only) and an indicator for whether the owner has a mortgage (homeowner model only).

We then regress our CZ housing price measure on log CZ size. Estimates are quite stable across the different measures, and point to an elasticity of housing costs with respect to log size that is around 0.2 or larger, substantially so when we limit to the 50 largest CZs.

[2] The quality adjustment models include controls for type of unit (with five classes of apartment building size), number of bedrooms and of total rooms, year of construction, indicators for utilities included in rent (for rents only), and an indicator for whether the owner has a mortgage (homeowner model only). The model for rents has an R-squared of 0.34, while that for home values has an R-squared of 0.50.

Appendix Figure 1. Wage changes for movers in and out of metropolitan areas (Glaeser and Mare, 2000)

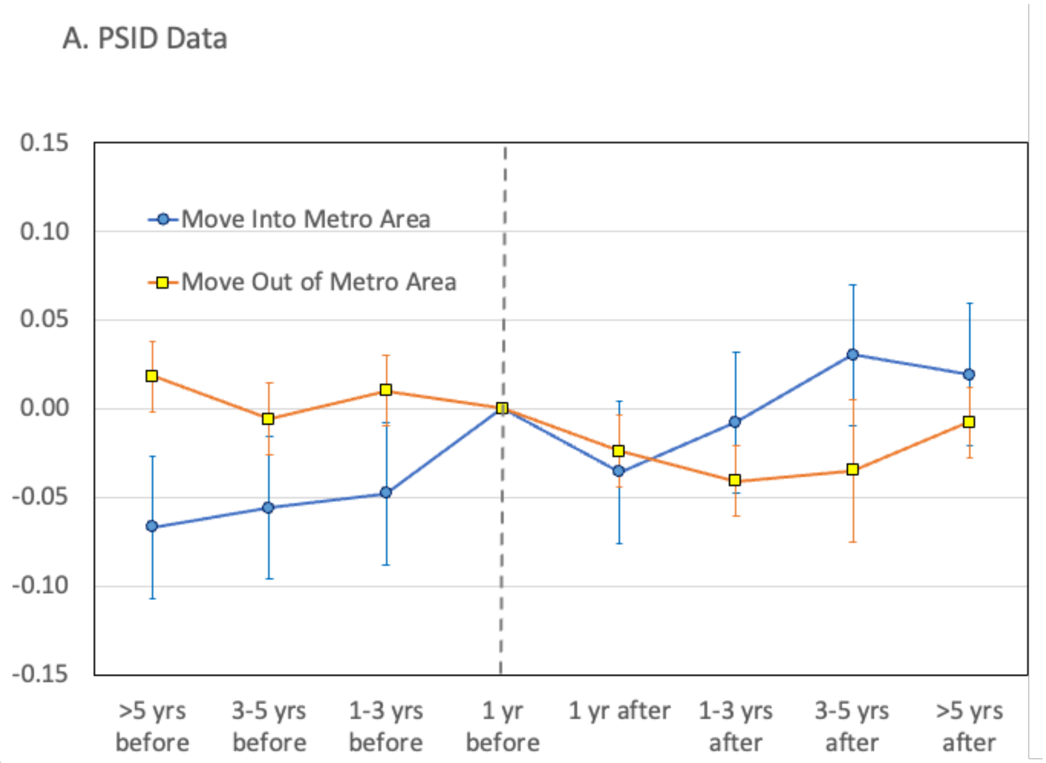

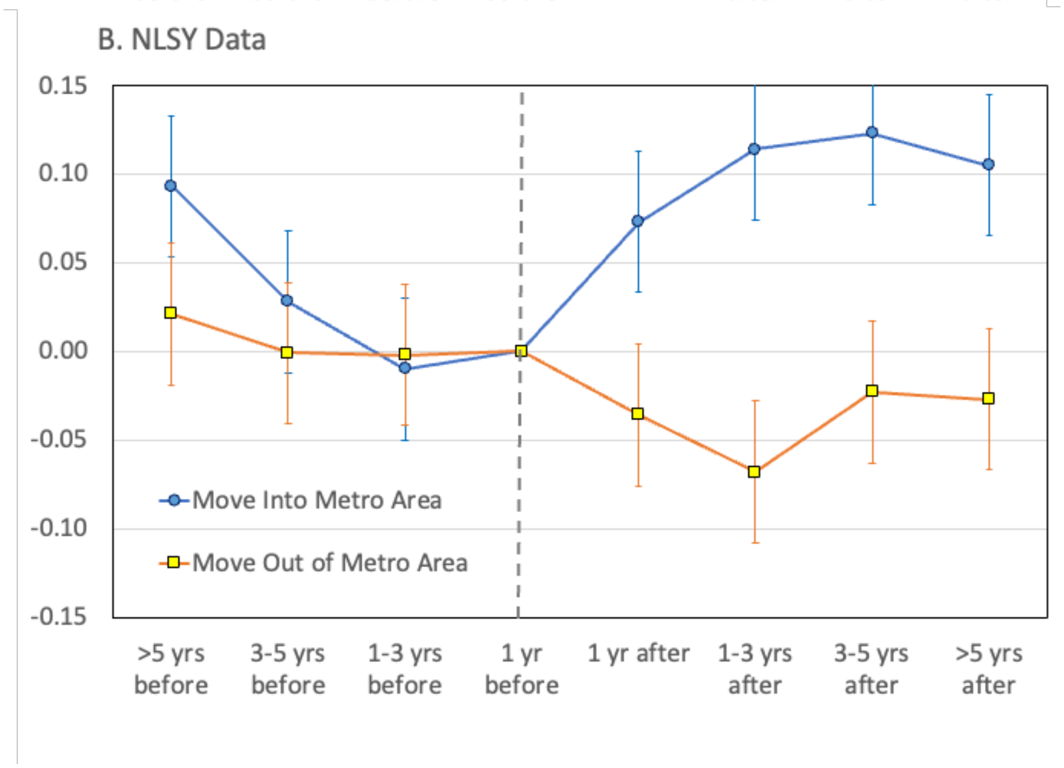

Notes: From Glaeser and Mare (2000, Table 5, columns 2 and 4). 95% confidence intervals shown with vertical bars.

Appendix Figure 2. Estimated CZ Wage Effects from Cross-Sectional Model versus Mean Log Wage in CZ, ACS data

Notes: CZ effects on vertical axis are fixed effects from our third cross-sectional model fit to the ACS, controlling for years of completed education (fixed effects), interacted with gender, immigrant status, and whether the immigrant has been in the U.S. for less than their potential experience + 5; experience (quartic, interacted with gender); race and ethnicity dummies (interacted with gender), year fixed effects; three immigrant source regions, interacted with years in the U.S.; field of degree (16 categories, interacted with gender, only for those with a BA or more); and detailed industry. Fixed effects are normalized to weighted mean zero across CZs. Scatterplot shows the 600 largest CZs. Overlayed lines show the 45-degree line (with intercept adjusted to equal zero for the average CZ) and a weighted regression line fit to all CZs (slope=0.69; robust standard error=0.02).

Appendix Figure 3. Relationship of Mean Log Wage and Cross-Sectional CZ Wage Effect to Log of CZ Size, ACS data

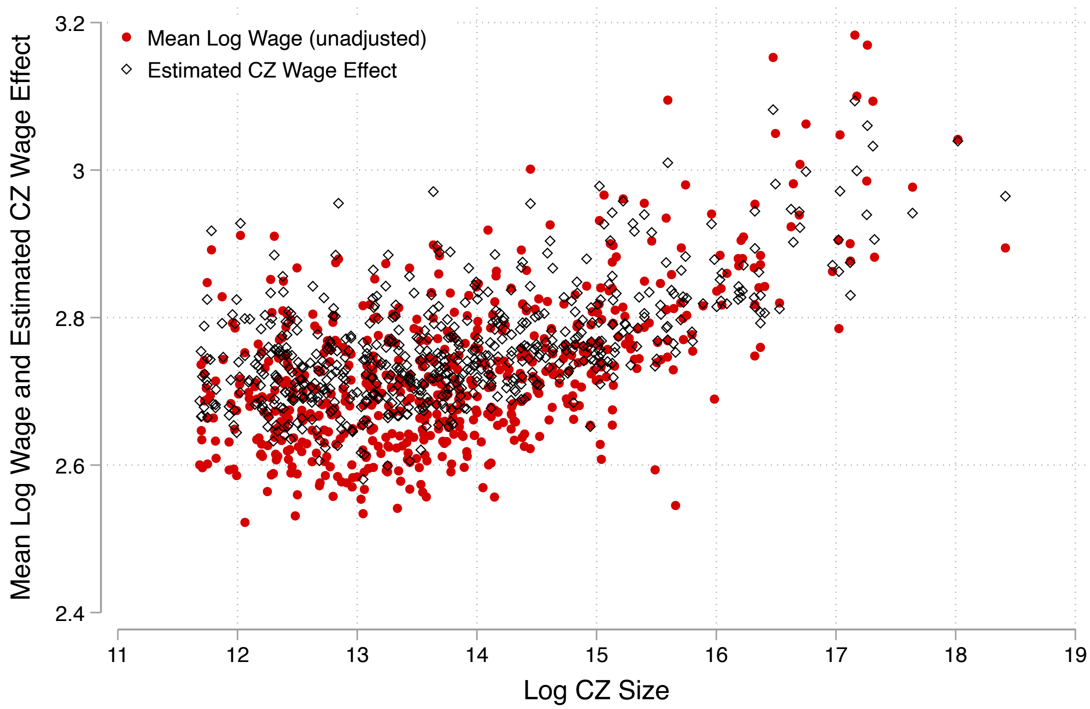

Notes: See notes to Appendix Figure 2 for description of estimated CZ wage effects. Points are shown for the 600 largest CZs. Wage effects are adjusted to have the same weighted mean as mean CZ wages.

Appendix Figure 4. Relation of Cross-Sectional CZ Wage Effects for High School and College Workers with Cross-Sectional CZ Wage Effect for All Workers, ACS data

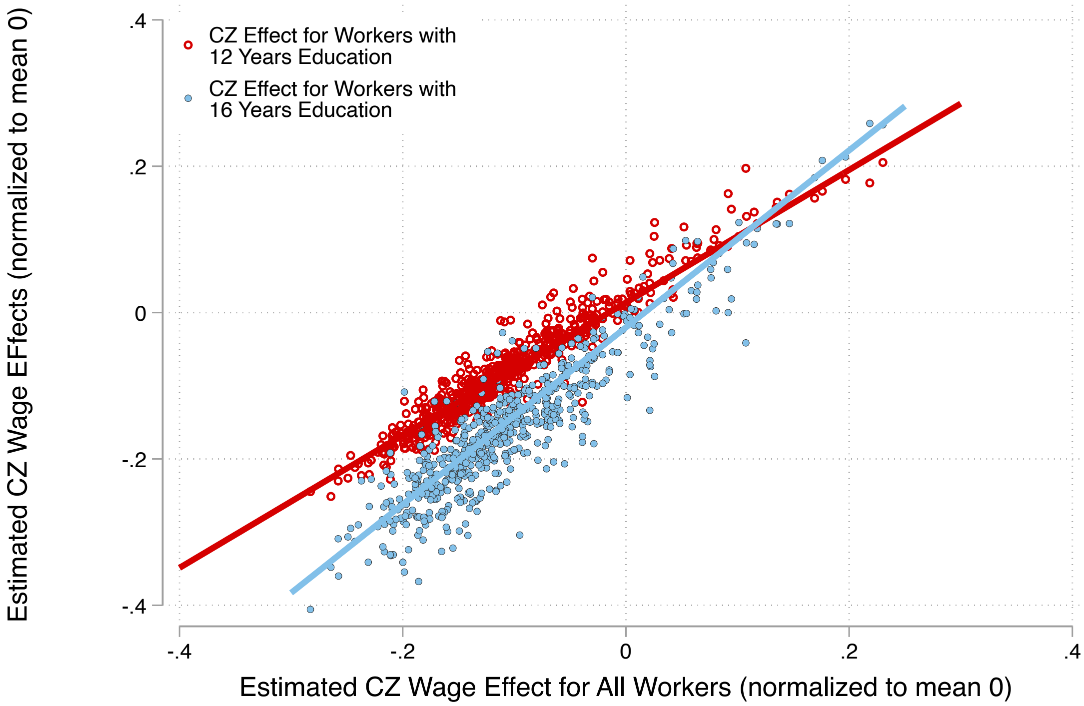

Notes: CZ effects on X-axis are from the model described in the notes to Appendix Figure 2. The model is then re-fit separately for workers with exactly 12 and exactly 16 years of education. CZ fixed effects from these models are plotted on the Y-axis for the 600 largest CZs. Overlayed lines are lines of best fit, fit to the full sample and weighted by the number of workers at each education level.

Appendix Figure 5. Mean AKM residuals by decile of person effect and CZ/industry effect

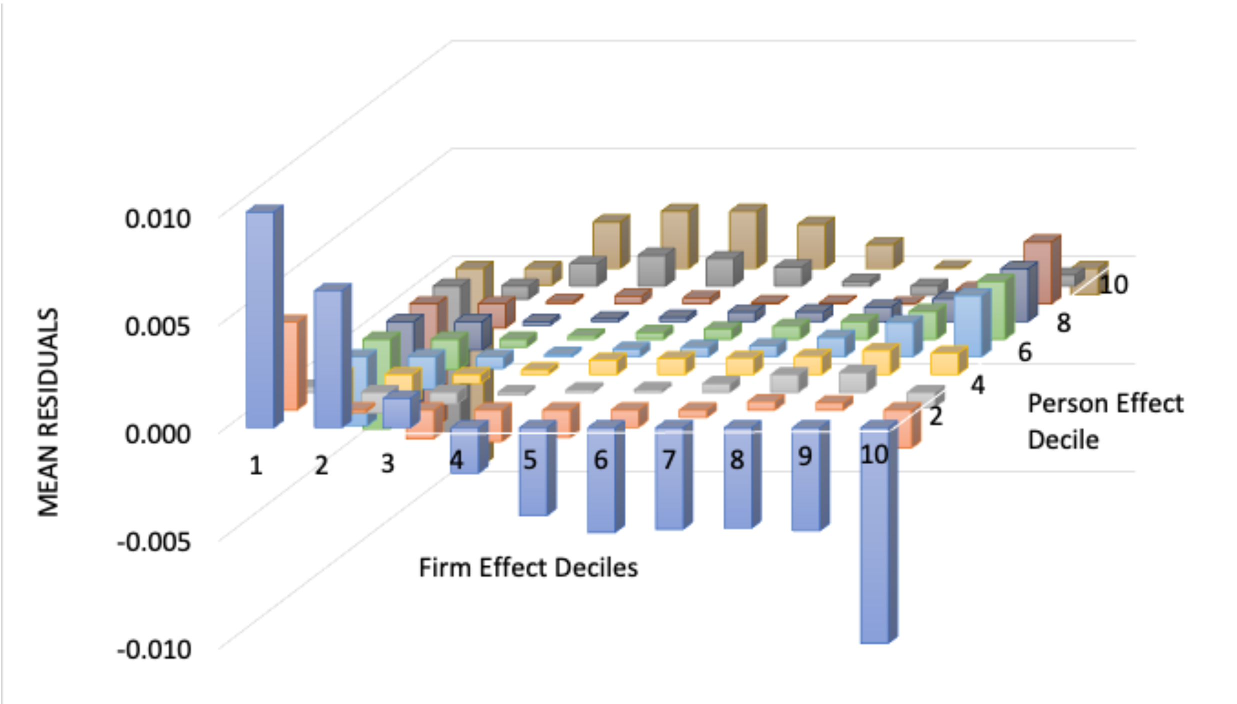

Note: Figure shows mean residuals from AKM specification by decile of the estimated person and firm effects.

Appendix Figure 6. Mean earnings residuals from AKM model before and after a change of CZs, by change in CZ earnings premium

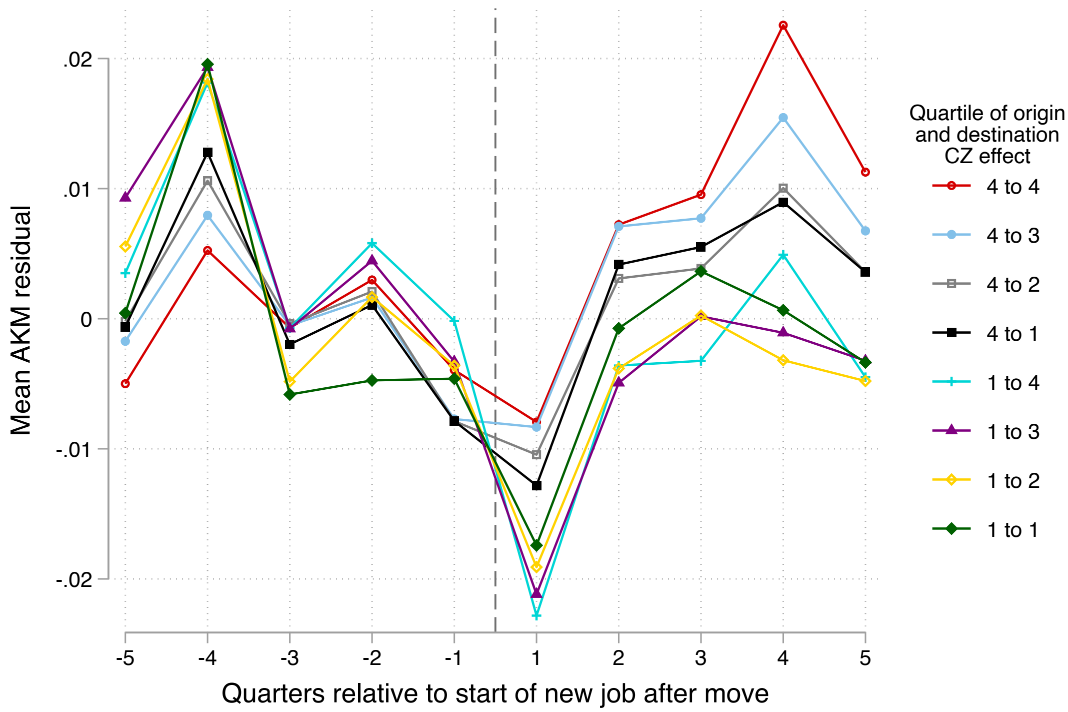

Notes: Figure shows event-time means of the AKM residual for workers who move between CZs, separately for different quartiles of the origin and destination AKM-based CZ earnings premiums. Quarter -1 represents the last full quarter (i.e., the second-to-last observed quarter) in the origin CZ, and quarter +1 represents the first full quarter (second observed quarter) in the destination CZ. We allow up to four quarters of non-employment between the two observed spells. Sample is limited to workers who move between CZs only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Appendix Figure 7. Mean earnings and mean residuals from AKM model before and after move across firms, by change in firm earnings premium

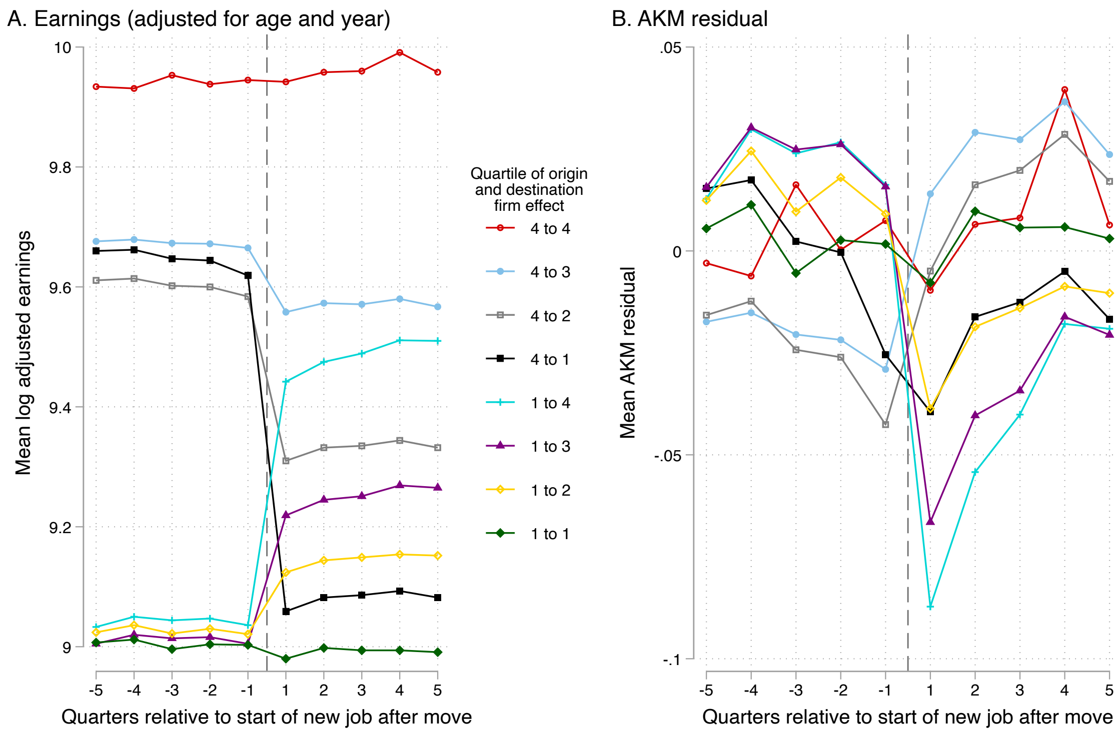

Notes: Figure shows event-time means for workers who move between establishments, separately for different quartiles of the origin and destination AKM-based establishment earnings premiums. Quarter -1 represents the last full quarter (i.e., the second-to-last observed quarter) in the origin establishment, and quarter +1 represents the first full quarter (second observed quarter) in the destination establishment. We allow up to four quarters of non-employment between the two observed spells. Sample is limited to workers who move between establishments only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Appendix Figure 8. Change in earnings and AKM residuals of firm movers from last pre-move quarter to first post-move quarter, by change in firm premium

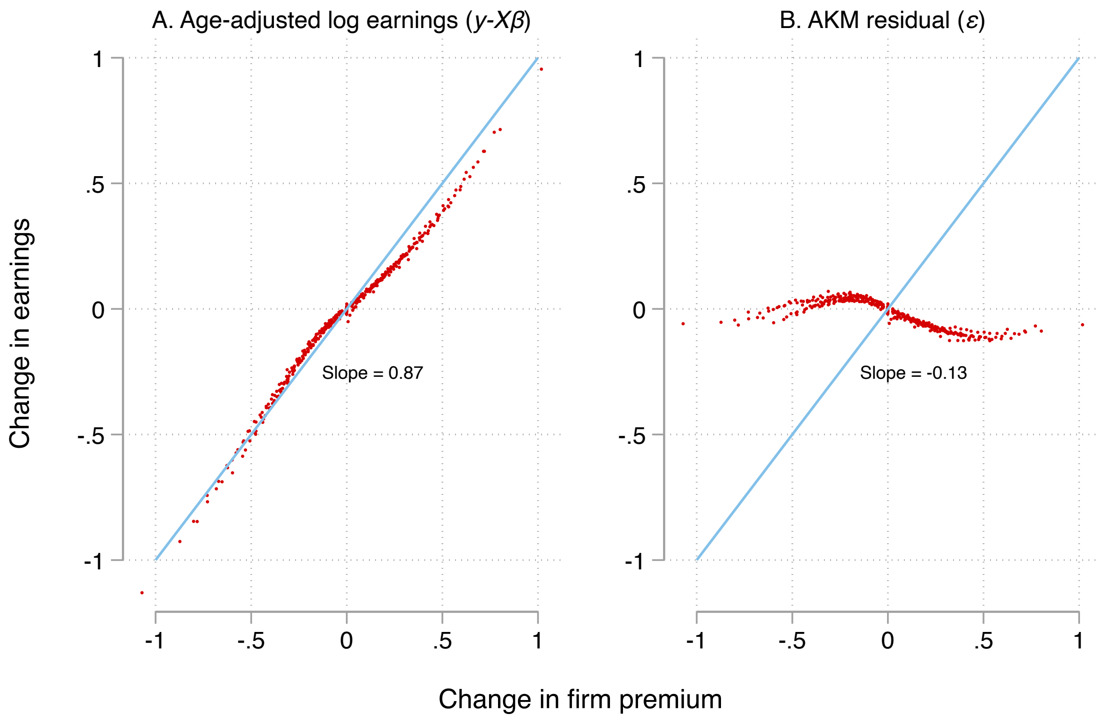

Notes: Firms are classified into 20 vingtiles based on their AKM-based establishment earnings premiums. We then classify firm movers based on the change in mean premium from their origin to their destination vingtile. Figure shows the average change in the indicated earnings components from the last full quarter in the origin firm to the first full quarter in the destination firm for each of the resulting 400 cells. Slopes correspond to the unweighted best linear fit line fit to these 400 points. Sample is limited to workers who move between establishments only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Appendix Figure 9. Comparison of earnings changes for “upward” and “downward” firm movers

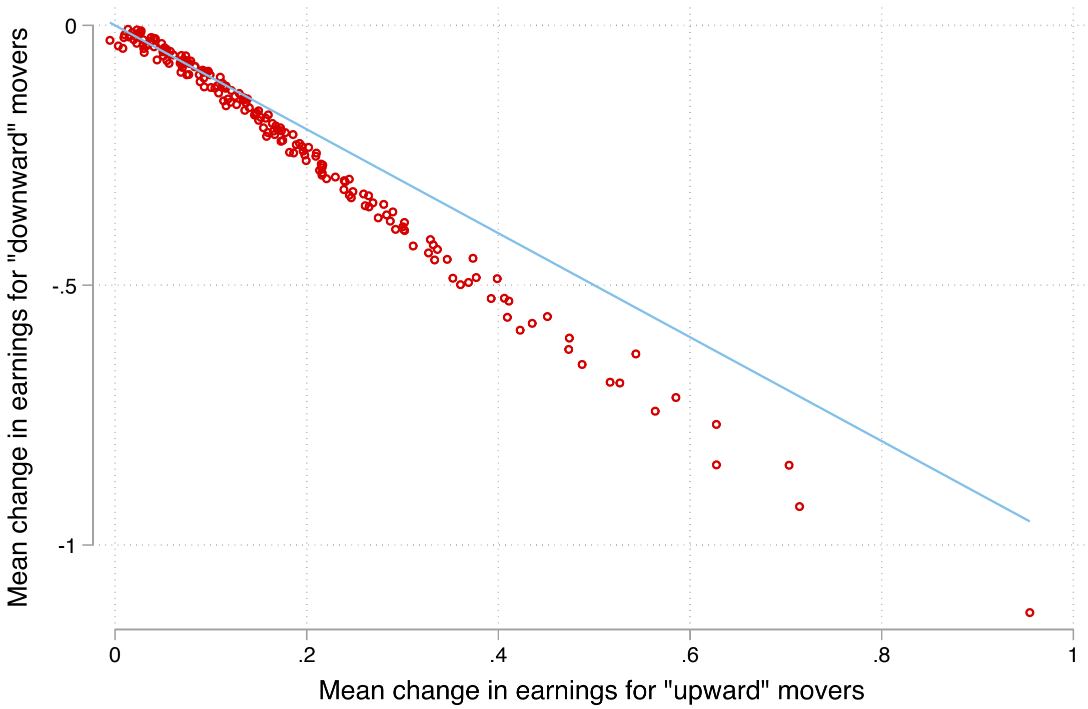

Note: Firms are classified into 20 vingtiles based on their AKM-based establishment earnings premiums. We identify all 190 origin-destination pairs of vingtiles where the destination has higher premiums than the origin and compute the average change in earnings from the last full pre-move quarter to the first full post-move quarter (excluding transition quarters) for this origin-destination combination. These values are plotted on the x-axis. On the y-axis, we plot the average change for movers going in the opposite direction, toward a lower-premium firm. The overlayed line represents what would be expected if average earnings changes were equal and opposite. Sample is limited to workers who move between establishments only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Appendix Figure 10. Comparison of earnings changes for “upward” and “downward” CZ movers

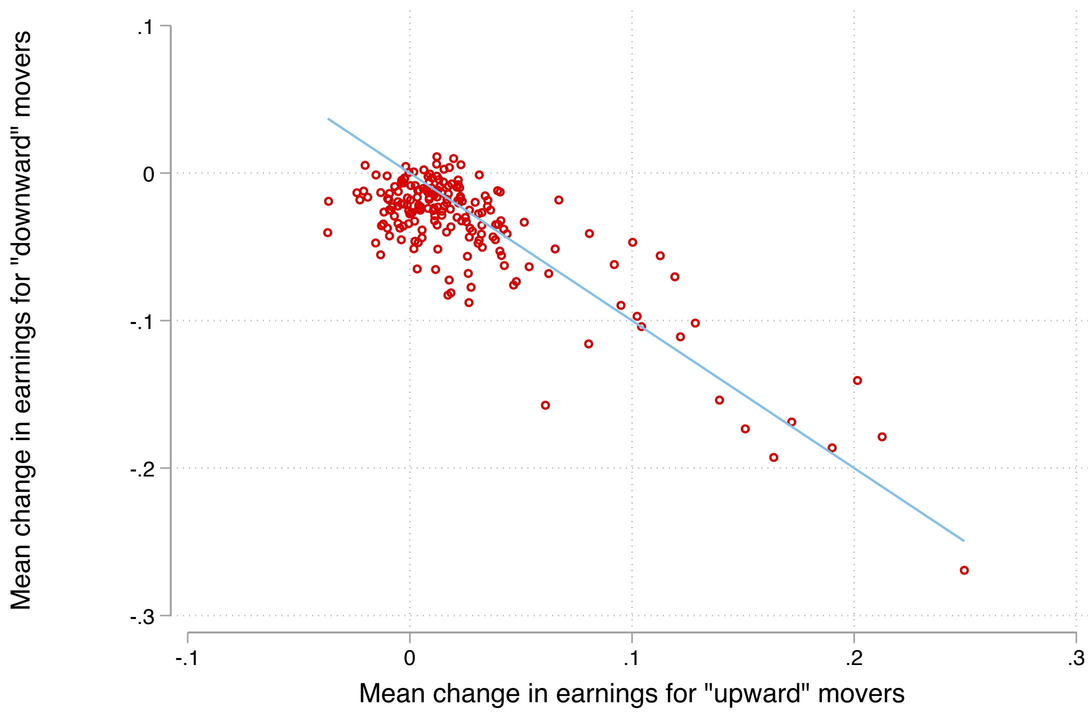

Note: CZs are classified into 20 vingtiles based on their AKM-based CZ earnings premiums. We identify all 190 origin-destination pairs of vingtiles where the destination is has higher premiums than the origin and compute the average change in earnings from the last full pre-move quarter to the first full post-move quarter (excluding transition quarters) for this origin-destination combination. These values are plotted on the x-axis. Note that a few are negative because the mean change in earnings has the opposite sign from the difference in estimated CZ premiums from our AKM model. On the y-axis, we plot the average change for movers going in the opposite direction, toward a lower-premium CZ. The overlayed line represents what would be expected if average earnings changes were equal and opposite. Sample is limited to workers who move between CZs only once in our sample and are observed with stable jobs at the same firm for at least five consecutive quarters (not including transition quarters) before and after the switch.

Appendix Table 1: CZ size effects on log hourly wages and log annual earnings, ACS data

	Unadjusted (1)	CZ Effect from Model 2 (2)	CZ Effect from Model 3 (3)
All Workers			
Log hourly wage	0.068	0.059	0.056
	(0.010)	(0.005)	(0.005)
	[0.497]	[0.606]	[0.604]
Log annual earnings	0.078	0.059	0.060
	(0.013)	(0.007)	(0.006)
	[0.489]	[0.540]	[0.589]
High School-educated Workers (12 years education)			
Log hourly wage	0.031	0.050	0.049
	(0.004)	(0.004)	(0.004)
	[0.306]	[0.547]	[0.546]
Log annual earnings	0.028	0.047	0.050
	(0.005)	(0.005)	(0.005)
	[0.214]	[0.427]	[0.511]
Some College or More Workers(>12 years education)			
Log hourly wage	0.082	0.070	0.065
	(0.009)	(0.005)	(0.005)
	[0.590]	[0.657]	[0.653]
Log annual earnings	0.092	0.072	0.069
	(0.013)	(0.007)	(0.006)
	[0.550]	[0.605]	[0.637]

Note: Table entries are estimated regression coefficient, robust standard error (in parentheses) and R-squared [in square brackets] from OLS regression of estimated CZ wage or earnings differential on log of CZ size. Models in column 1 are fit to unadjusted mean log hourly wages or mean log annual earnings. Models in columns 2 and 3 are fit to adjusted CZ earnings differentials derived from models described in text. All models are fit by weighted least squares using as a weight the total number of worker observations for the CZ in the 2010-2018 ACS sample.

Appendix Table 2: Models for CZ employment rates and average hours, ACS data

	Both Genders			Males only			Females only		
		Regressed on mean log wage in CZ			Regressed on mean log wage in CZ			Regressed on mean log wage in CZ	
	Regressed on log(size)	Unadjusted	Adjusted	Regressed on log(size)	Unadjusted	Adjusted	Regressed on log(size)	Unadjusted	Adjusted
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable = CZ-mean of:									
Employed last year	0.005	0.15	0.12	0.010	0.18	0.17	0.001	0.13	0.07
	(0.003)	(0.02)	(0.03)	(0.003)	(0.02)	(0.03)	(0.004)	(0.02)	(0.04)
Employed last year (adjusted)	0.000	0.07	0.04	0.001	0.06	0.04	-0.001	0.08	0.04
	(0.002)	(0.02)	(0.03)	(0.002)	(0.02)	(0.03)	(0.002)	(0.01)	(0.03)
Hours last year	11.42	337	283	15.64	359	346	8.97	324	229
	(8.28)	(49)	(83)	(8.22)	(53)	(86)	(8.98)	(50)	(87)
Hours last year (adjusted)	-3.84	81	35	-7.44	35	-7	-0.01	127	79
	(5.37)	(40)	(63)	(6.12)	(49)	(75)	(4.88)	(34)	(54)

Notes: Table entries are regression coefficients from univariate regressions of dependent variable shown in row heading on independent variable indicated in column heading. Hours last year includes zeros for non-workers. All models are estimated using data for 691 CZ's and characteristics estimated using 2010-2018 ACS files. "Adjusted" means that the dependent variable or the independent variable (or both) is estimated from a model with CZ dummies and a set of individual characteristics. All regression models are weighted by the (weighted) count of working-age people in the CZ in the ACS files. Robust standard errors in

Appendix Table 3. Predictions of CZ wage premiums based on public data (top 50 CZs)

Rank	CZ number	Largest city in CZ	Approx. premium	Rank	CZ number	Largest city in CZ	Approx. premium
1	38300	Los Angeles, CA	0.050	26	16300	Pittsburgh, PA	-0.053
2	19400	New York, NY	0.098	27	33000	Fort Worth, TX	0.010
3	24300	Chicago, IL	0.046	28	7400	Orlando, FL	-0.065
4	11304	Washington DC	0.154	29	38801	Portland, OR	0.014
5	19600	Newark, NJ	0.113	30	18000	Buffalo, NY	-0.040
6	32000	Houston, TX	0.046	31	31301	San Antonio, TX	-0.026
7	19700	Philadelphia, PA	0.046	32	37901	Las Vegas, NV	0.016
8	20500	Boston, MA	0.093	33	31201	Austin, TX	0.022
9	37800	San Francisco, CA	0.169	34	12701	Cincinnati, OH	-0.020
10	9100	Atlanta, GA	0.006	35	29502	Kansas City, MO	-0.021
11	11600	Detroit, MI	-0.031	36	1701	Raleigh, NC	-0.011
12	39400	Seattle, WA	0.074	37	15900	Columbus, OH	-0.024
13	33100	Dallas, TX	0.039	38	900	Charlotte, NC	-0.010
14	7000	Miami, FL	-0.035	39	14200	Indianapolis, IN	-0.028
15	35001	Phoenix, AZ	0.003	40	36100	Salt Lake City, UT	-0.022
16	21501	Minneapolis, MN	0.036	41	24100	Milwaukee, WI	-0.010
17	20901	Bridgeport, CT	0.090	42	5600	Nashville-Davidson, TN	-0.037
18	28900	Denver, CO	0.036	43	7100	West Palm Beach, FL	-0.020
19	38000	San Diego, CA	0.055	44	20401	Providence, RI	0.038
20	37400	Sacramento, CA	0.051	45	7600	Jacksonville, FL	-0.031
21	11302	Baltimore, MD	0.094	46	37200	Fresno, CA	0.008
22	6700	Tampa, FL	-0.046	47	12200	Grand Rapids, MI	-0.066
23	37500	San Jose, CA	0.183	48	20600	Manchester, NH	0.017
24	15200	Cleveland, OH	-0.045	49	33803	Oklahoma City, OK	-0.057
25	24701	St. Louis, MO	-0.025	50	2000	Virginia Beach, VA	-0.020

Notes: CZ premiums are best linear predictors of non-disclosed LEHD-based premiums from public ACS data, using as predictors the estimated premium from our first cross-sectional model, the log size of the CZ, and the share of workers in the CZ with some college or more. Premiums are normalized to weighted mean zero across CZs.

Appendix Table 4: Comparison of CZ-Level Variance Decompositions for Main Sample, No-Uncertainty Sample, and Non-Mover Sample

	Main sample		Non-mover sample		No-uncertainty sample	
	Std. Dev. or Correlation	Var. Share	Std. Dev. or Correlation	Var. Share	Std. Dev. or Correlation	Var. Share
	(1)	(2)	(3)	(4)	(5)	(6)
Log earnings or mean log earnings	0.147	100%	0.140	100%	0.145	100%
Variance components (std. deviations in odd-number columns, variance shares in even-numbered columns)						
Person effects	0.081	30.3%	0.076	29.2%	0.081	31.0%
Firm effects	0.079	29.3%	0.077	30.4%	0.077	28.3%
Covariate index (Xβ)	0.006	0.2%	0.006	0.2%	0.006	0.2%
Covariance components (correlations in odd-number columns, variance shares in even-numbered columns)						
Person & firm	0.642	38.2%	0.617	36.8%	0.643	38.1%
Person & covariate index	0.228	1.0%	0.337	1.5%	0.259	1.2%
Firm & covariate index	0.245	1.0%	0.406	1.9%	0.239	1.1%
Sample size (billions of person-quarter observations)	2.52		1.84		1.90	

Notes: Table shows variance decompositions of CZ-level mean earnings based on equation (6). Entries in odd-numbered columns "variance components" are standard deviations of the earnings components indicated in row headings; for "covariance components" they are the estimated correlations of the indicated components. Entries in even-numbered columns are variance shares explained by the variance or covariance components.

Appendix Table 5: Summary of AKM model estimates by education group

	Low education sample				High education sample			
	Person-quarter level		CZ level		Person-quarter level		CZ level	
	Std. Dev. or Correlation	Var. Share	Std. Dev. or Correlation	Var. Share	Std. Dev. or Correlation	Var. Share	Std. Dev. or Correlation	Var. Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log earnings or mean log earnings	0.514	100%	0.097	100%	0.678	100%	0.178	100%
<u>Variance components (std. deviations in odd-number columns, variance shares in even-numbered columns)</u>								
Person effects	0.409	63.5%	0.055	32.3%	0.559	67.8%	0.101	32.2%
Firm effects	0.212	17.0%	0.059	36.8%	0.223	10.8%	0.087	24.1%
Covariate index (Xβ)	0.127	6.1%	0.007	0.6%	0.130	3.7%	0.005	0.1%
Residual	0.192	14.0%	0.001	0.0%	0.244	13.0%	0.000	0.0%
<u>Covariance components (correlations in odd-number columns, variance shares in even-numbered columns)</u>								
Person & firm	0.112	7.3%	0.433	29.9%	0.167	9.0%	0.753	41.9%
Person & covariate index	-0.177	-7.0%	-0.247	-2.1%	-0.139	-4.4%	0.318	1.1%
Firm & covariate index	-0.017	-0.4%	0.226	2.1%	0.015	0.2%	0.255	0.7%

Notes: Table shows variance decompositions based on equation (4). Columns 1-2 and 5-6 pertain to the variance of individual quarterly earnings. Columns 3-4 and 7-8 pertain to the variance of mean earnings by CZ. Entries in odd-numbered columns for "variance components" are standard deviations of the earnings components indicated in row headings; for "covariance components" they are the estimated correlations of the indicated components. Entries in even-numbered columns are variance shares explained by the variance or covariance components.

Appendix Table 6: Variance Decomposition of CZ Wage Gap between High- and Low-education Workers

	Variance Decomposition of Wage Gap	
	Std. Dev. or Correlation (1)	Var. Share (2)
Wage gap (high- versus low-education workers)	0.107	100%
Components of Wage Gap (column 1 = std. dev.)		
Difference in mean person effects	0.074	49%
Difference in covariate indexes	0.006	0%
Difference in mean CZ wage effect:	0.038	13%
Of which:		
Relative wage premium	0.014	2%
Composition	0.028	7%
Interaction	0.002	0%
Covariance Terms (column 1 = correlation of terms)		
Cov(Person effects, cov. index)	-0.251	-2%
Cov(Person effects, CZ effect)	0.831	42%
Of which:		
Cov(Person effects, relative wage premium)	0.785	14%
Cov(Person effects, composition)	0.785	28%
Cov(Person effects, interaction)	-0.311	-1%
Sum of all other covariance terms	--	3%

Notes: Table shows variance decomposition of the difference in mean log earnings of college and non-college workers in a CZ, using equations (12) and (13). Analysis is weighted by CZ size.

Appendix Table 7: Elasticities of Housing Values and Rents with respect to CZ Size

	All CZ's (1)	**Largest 50 CZ's** (2)
Housing Prices (log of home value for owners)		
Unadjusted	0.25	0.38
	(0.01)	(0.08)
Quality Adjusted	0.22	0.42
	(0.01)	(0.08)
Monthly Rent (log of rent for renters)		
Unadjusted	0.17	0.19
	(0.01)	(0.04)
Quality Adjusted	0.18	0.23
	(0.01)	(0.04)

Note: Table entries are regression coefficients (and standard errors) from weighted OLS regressions of CZ-average housing price measure in row heading on constant and log of number of workers in CZ. Regressions are weighted by number of workers in CZ. Sample in column 1 is set of 678 CZ's in 2018 5-year ACS with non-missing data. Sample in column 2 is 50 largest CZ's, ranked by number of workers. Quality adjusted values and rents derive from regressions on indicator for type of housing unit, number of bedrooms, log of total number of rooms, year of construction, and indicator for mortgage (for home values) or set of indicators for inclusion of utility costs (for rents).