Federal Reserve Bank of Cleveland Working Paper Series

Labor Market Tightness and Minority Earnings Gaps

Vittoria Dicandia and Victor Hernandez Martinez

Working Paper No. 26-34

November 2026

Suggested citation: Dicandia, Vittoria and Victor Hernandez Martinez. 2026. "Labor Market Tightness and Minority Earnings Gaps." Working Paper No. 26-34. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202634.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

See more working papers at: www.clevelandfed.org/research. Subscribe to email alerts to be notified when a new working paper is posted at: https://www.clevelandfed.org/subscriptions.

This work is licensed under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International. To view a copy of this license, visit https://creativecommons.org/licenses/by-nc-nd/4.0/

Labor Market Tightness and Minority Earnings Gaps

Vittoria Dicandia [†] Victor Hernandez Martinez[‡]

September 2026

Abstract

We study how labor-market tightness affects minority earnings gaps using administrative data for US metropolitan areas from 2010 to 2025. After accounting for national shocks, local seasonality, and local trends, a 1-percentage-point unemployment decline reduces the white-minus-minority earnings gap by about 0.2 percentage points for all jobs and 0.3 percentage points for new stable hires. These responses strengthen as slack declines, reaching 0.4 and 0.6 percentage points, respectively, in hot markets. Neither observed worker and vacancy composition nor concurrent policy changes account for these effects. Increased mobility and upgrading across broad job categories explain little of the narrowing, although hiring distributions become more similar between the two groups. Where recruiting is concentrated among fewer employers, the earnings-gap response varies less with slack. Together, the findings point toward relative earnings improvements within labor-market segments and employer competition as a mechanism through which tightness narrows minority earnings gaps.

Keywords: Labor-market tightness; minority earnings gaps; job mobility; employer competition; monopsony

JEL Codes: J63, J62, J31, J71, J42, E24

We thank Peter Kuhn, Ioana Marinescu, Roberto Pinheiro, Kevin Rinz, and Benjamin Schoefer for their helpful advice. Nilay Bhadra, James Hamlette, and Geena Panzitta provided excellent research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors, the Federal Reserve Bank of Cleveland, or the Federal Reserve System.

† Federal Reserve Bank of Cleveland, 1455 E 6th, Cleveland, OH, USA. email: vittoria.dicandia@clev.frb.org

‡ Federal Reserve Bank of Cleveland, 1455 E 6th, Cleveland, OH, USA. email: victor.hernandezmartinez@clev.frb.org

1 Introduction

Tighter labor markets raise employment and the wages employers must pay to attract and retain workers. These gains are unevenly distributed: employment responds more strongly to labor-market conditions for less-skilled and minority workers, on both the extensive and the intensive margin (Bound and Holzer 2000; Cajner et al. 2017; Hines et al. 2001; Hoynes 2000). Whether tightening also narrows the gap in *pay* between minority and white workers is far less settled. Economic theory provides no unambiguous prediction, because tightness operates through several offsetting margins at once: the composition of employment, the segments in which labor demand expands, and the pay of observationally similar workers (Black 1995; Solon et al. 1994).

We study how labor-market tightness affects the earnings gap between white workers and Black and Hispanic workers using administrative MSA-level data from 2010 to 2025. Our local-market approach allows us to study the earnings effects of tightness after accounting for national shocks, local seasonality, and gradual demographic and economic changes. Tighter labor markets narrow minority earnings gaps, with larger reductions among new hires. These effects strengthen as labor-market slack declines, particularly among new hires. Observed composition and policy changes, increased mobility, and upgrading across broad job categories explain little of this compression, though the job types the two groups enter do become more similar. Employer competition, however, matters: recruiting concentration limits the additional compression that tightening produces as markets heat up. Together, our findings argue against labor-market segmentation and reallocation, and point to employer competition as a relevant channel through which tightness affects earnings gaps.

We combine quarterly earnings by group from the Census Bureau's Quarterly Workforce Indicators (QWI), built from aggregated worker–employer records, with local labor-market conditions for up to 374 MSAs. These data identify minority earnings gaps market by market, whereas much of the previous evidence on tightness and racial gaps rests on national or state aggregates. We regress the minority earnings gap, measured as 100 times the log difference between white and pooled Black and Hispanic mean monthly earnings, on tightness across MSA-quarters. We measure tightness by the negative of the local unemployment rate, because we show that local job postings carry no additional information about the gap once aggregate movements are absorbed. Our specification absorbs the variation most likely to confound the estimates: conditions common to all markets in a quarter, each market's recurring seasonal pattern, and its own linear trend.

We therefore identify the response from a market's deviation from its own seasonal path, net of national developments and of gradual local demographic and economic changes.

A 1-percentage-point decline in unemployment reduces the minority earnings gap by about 0.2 percentage points for all jobs and 0.3 percentage points for new stable hires. These magnitudes are equivalent to the effect of a one-dollar increase in the real minimum wage estimated in the same specification. As labor-market slack declines, the effects of tightness strengthen and the difference between all jobs and new stable hires widens: in cold markets the same unemployment decline reduces the gap by approximately 0.2 percentage points for both, whereas in hot markets the reductions reach 0.4 percentage points for all jobs and 0.6 percentage points for new stable hires. Alternative measures of tightness deliver the same pattern, and estimates within six macroeconomic windows since 2002 show that our results are not driven by a single episode.

Our results do not reflect changes in the workers or jobs being compared, or policies that coincide with tightening. Adding MSA-level workforce-composition measures, including differences in education, age, and other earnings-related characteristics, leaves the estimates unchanged. Selective geographic mobility does not account for the compression either: workers entering and leaving metropolitan areas contribute little to changes in the minority earnings gap. On the demand side, we construct minority- and white-specific measures of tightness and show that stronger demand for the kinds of jobs minorities disproportionately hold does not drive the effect of aggregate tightness on the earnings gap. Finally, the estimates remain unchanged after accounting for minimum wage changes and job-posting measures of employers' diversity, equity, and inclusion (DEI) engagement. Taken together, these exercises provide little support for observed compositional or concurrent policy changes as explanations for the earnings compression.

Job mobility and broad job-type upgrading do not account for the effect of tightness on the earnings gap. Using the Census Bureau's Job-to-Job Flows, we find that hiring from nonemployment and moves between employers rise more strongly for minorities than for whites. Accounting for differences between the groups in transition rates, however, leaves the earnings response essentially unchanged for all jobs and only partly reduces it for new stable hires. We also find limited evidence of systematically greater minority upgrading across broad job categories. Nevertheless, the distributions of minority and white hires across industry–firm-size job types converge as markets tighten. Therefore, our findings suggest that tightening narrows minority earnings gaps within job categories, pointing toward the importance of employer-specific dimensions in

explaining these effects.

Employer competition shapes how strongly tightening compresses the earnings gap. Employer competition can improve earnings because workers may receive more and better offers, or because their current employers may raise pay to retain them. We construct a proxy for employer market power using the concentration of local recruiting across employers, measured from job postings within detailed occupations (Azar et al. 2020), and interact it with tightness. Where recruiting is concentrated, the earnings-gap response to tightness varies less with slack. The difference is clearest in hot labor markets: a 1-percentage-point decline in unemployment narrows the gap for all jobs by about 0.4 percentage points where recruiting is concentrated, against 0.6 percentage points where it is dispersed. These results point toward employer competition as an important channel through which improvements in labor-market conditions affect earnings gaps.

Our findings speak to three families of explanations for minority earnings gaps: the kinds of jobs workers hold, how employers evaluate and pay workers, and the set of jobs and employers workers can effectively reach. The composition, occupational-demand, and destination results restrict the first, since neither who is employed nor reallocation across broad job categories accounts for the compression. On the other hand, our results cannot differentiate between explanations based on discrimination and those based on unequal access to job opportunities. The two are observationally equivalent in our setting even though the disadvantage they describe differs: tighter markets improve workers' outside options, and better outside options can reduce the earnings consequences of discrimination and of unequal access alike, without any change in employers' preferences or in workers' networks. What the evidence does indicate is that competition among employers changes the earnings consequences of these disadvantages, whatever their origin.

This paper contributes to several strands of the literature. Research on strong labor markets finds particularly large employment and hours gains for disadvantaged workers (Hines et al. 2001). Related work on business-cycle fluctuations finds larger cyclical employment responses among disadvantaged workers, while differences in earnings responses are more heterogeneous (Hoynes 2000). More recently, Autor et al. (2024) link post-pandemic wage compression at the bottom of the distribution to stronger employer competition and note a contemporaneous narrowing of racial wage gaps. Our work provides a systematic analysis of how local labor-market tightness shapes minority earnings gaps and of the labor-market margins underlying that relationship. Aaronson et al. (2019) find additional high-pressure employment gains for disadvantaged groups, but lit-

tle evidence of improvements in relative wages or earnings. Using granular administrative data across local labor markets, we not only provide more robust evidence on the earnings margin but also convincingly document that minority earnings-gap compression becomes stronger as slack declines. We also document a stronger response among new stable hires and its stronger variation with the degree of labor-market slack, which extends research on the cyclicality of wages in newly formed employment relationships (Gertler et al. 2020; Martins et al. 2012) to earnings inequality between whites and minorities.

We also contribute to the literature on the cyclical job ladder (Haltiwanger et al. 2018; Moscarini and Postel-Vinay 2018) and racial earnings inequality by examining whether differences in job mobility responses to tightness between whites and minorities account for earnings convergence. Racial differences in hiring and job loss drive employment disparities over the cycle (Cajner et al. 2017; Couch and Fairlie 2010). Additionally, the returns to mobility themselves differ by race. For instance, higher job churn among Black men and closer monitoring of Black workers reentering employment widen the earnings gap (Cavounidis et al. 2024; Daly et al. 2020). Our results show that minority-relative job mobility rises with tightness, yet accounting for these differences leaves most of the earnings compression intact, consistent with our limited evidence of minority-relative upgrading across broad job categories. Because the pay ranking of a broad category is itself an incomplete measure of job quality (Sorkin 2018), our findings do not rule out meaningful advancement along dimensions those categories do not capture. Indeed, our results on hiring distributions show that they become more similar between whites and minorities with tighter markets. Overall, our findings suggest that tightening narrows minority earnings gaps within broad job categories, potentially through changes in hours or employment continuity, movement toward better-paying employers (Gerard et al. 2021; Sorkin 2025), or earnings gains within existing employment relationships.

Finally, we connect earnings inequality between whites and minorities to the interaction between employer competition and labor scarcity in two ways. First, our results highlight the potential role of outside opportunities in strengthening workers' bargaining positions even when they do not change employers, since transition frequency and upgrading across broad job categories appear to explain little of the earnings compression. This connects our work to Lang et al. (2005), who show how discriminatory hiring can weaken minority workers' alternatives and lower their earnings, and to Wursten and Reich (2023), who link racial wage convergence following minimum-wage

increases to improved access to alternative employers. More broadly, our results connect to the literature on how outside offers and worker mobility shape employer wage-setting power (Cahuc et al. 2006). Second, we speak directly to the role of employer competition by showing that when recruiting is more concentrated across employers, minority earnings-gap compression is less responsive to an increase in tightness. This finding extends the literature on how barriers to mobility and recruiting concentration limit the competitive pressure exerted by outside opportunities. In particular, it builds on Johnson et al. (2025), who show that restrictions on mobility can limit workers' gains from tight labor markets and widen racial earnings gaps, and on Burya et al. (2023), who show that recruiting concentration weakens the response of advertised wage growth to unemployment. Together, these contributions highlight the importance of employer competition in determining labor-market differences between minority and white workers.

2 Tightness and Minority Earnings Gaps

Tighter labor markets need not benefit all workers equally (Aaronson et al. 2019). Workers can gain by finding employment or increasing their earnings (via higher wages in their current jobs, greater job stability, more hours, or transitions to better-paying or higher-quality jobs), but the relative importance of these channels differs across workers. For example, Hines et al. (2001) document especially strong employment responses among less-skilled workers during expansions, while Autor et al. (2024) connect stronger employer competition to larger wage gains in lower-paying jobs among those working.

This heterogeneity makes a comparison between minority and white workers relevant because the two groups differ in their characteristics and positions in the labor market (Altonji and Blank 1999). Moreover, even when workers are observationally similar, employment circumstances may still generate group-specific responses to tightening. For example, Daly et al. (2020) show that Black and white workers with similar observed characteristics follow different earnings trajectories, with slower early-career wage growth among Black workers. Beyond whether minorities gain more from tightening, we also ask whether their relative earnings gains vary with the degree of labor-market slack. This question builds on evidence of additional employment benefits for some disadvantaged groups in high-pressure economies (Aaronson et al. 2019).

The literature motivates examining several channels through which tightening may affect minority earnings gaps. First, hiring and job loss contribute differently to racial employment gaps over

the cycle (Couch and Fairlie 2010), so we examine entry, separations, and stable employment separately. Second, the earnings consequences of changing employers depend not just on how often workers move but also on the types of jobs they move to. For instance, Haltiwanger et al. (2018) document procyclical movement to higher-paying firms. Therefore, we analyze destination pay and earnings around moves alongside transition frequency. Third, we study earnings responses to tightening separately for new hires and the broader stock of jobs because the literature shows that changes in the quality of newly formed matches can make new-hire earnings more responsive to labor-market conditions (Gertler et al. 2020). Finally, we investigate whether recruiting concentration shapes minorities' earnings gains relative to whites as markets tighten, since equally strong labor demand may generate less competition when recruiting is concentrated among a few firms. Such competition can benefit workers even without a move, as outside offers induce employers to raise pay to retain them (Cahuc et al. 2006).

The relevance of tightness for minority earnings gaps does not depend on the precise origins of these gaps. As we discuss in Section 10, segmented markets, discrimination, and information and network barriers provide different, and potentially complementary, channels through which tightness can shape the evolution of these gaps. Minority earnings gaps can narrow if minorities and whites belong to different segments of the labor market and labor demand increases disproportionately in minority-concentrated segments. Discrimination can restrict access to jobs and affect pay through different mechanisms (Lang and Lehmann 2012), but tightening generally increases employers' costs of maintaining their racial preferences as workers' alternatives improve. Personal networks or information barriers can create earnings disadvantages for workers with poorer access to well-paying jobs, but tighter markets can weaken those disadvantages by inducing employers to search more broadly, making access to job opportunities less dependent on personal contacts. Therefore, understanding how tightness changes minority earnings gaps does not require first settling their origins.

3 Data

We use the Census Bureau's Quarterly Workforce Indicators (QWI) to measure local earnings, employment, and hiring by race. The QWI are produced by the Longitudinal Employer–Household Dynamics (LEHD) program from unemployment-insurance wage records linked to information on workers and employers (US Census Bureau 2026a). They provide consistent measures of earn-

ings and employment flows for small geographic areas that are difficult to study with household surveys alone. Our main panel covers metropolitan statistical areas (MSAs) from the third quarter of 2010 through the second quarter of 2025, with up to 387 MSAs observed over this period. The data cover private-sector and state and local government jobs, accounting for approximately 96 percent of US wage-and-salary jobs (Jones-Ruiz 2020).[1] Appendix C.1 provides the construction details.

We study earnings for two job populations. The first consists of jobs held at the beginning of the quarter, including those that end before the quarter is over. The second consists of new stable hires: workers who joined an employer in the previous quarter and whose jobs continue through the reference quarter, excluding recent recalls. In both cases, the outcome is average monthly earnings.[2] New-stable-hire earnings therefore refer to the full quarter following entry, not the initial hiring quarter. These are earnings measures, not hourly wage rates: they reflect pay, hours, and time employed, and a stable job need not be full-time.

We compare non-Hispanic white workers with a pooled group of non-Hispanic Black workers and Hispanic workers of any race, which we collectively refer to as minorities. For each job population, MSA, and quarter, we first combine Black and Hispanic earnings using the corresponding job counts as weights. We then define the minority earnings gap as 100 times the difference between the natural log of white mean earnings and the natural log of this pooled minority mean. The outcome therefore describes the unadjusted earnings difference between the groups employed locally. A positive gap denotes higher white earnings, and a decline denotes earnings-gap compression.

The minority earnings gaps in the QWI data are consistent with those documented in previous work using household surveys (Bayer and Charles 2018; Daly et al. 2017). Appendix Figure C.1 shows that QWI gaps exhibit differences across states similar to those in the Current Population Survey (CPS) and American Community Survey (ACS), including after common annual movements are removed. We use states for this comparison because it is the most granular level of geography consistently identified across all three sources.

To limit the influence of earnings averages based on small groups of workers, we exclude MSAs with fewer than 100 new stable hires among either white workers or pooled minorities in any quarter of 2010Q1–2025Q2. This removes 13 of the 387 MSAs available in the QWI, leaving up to

[1] Federal employment and self-employment are excluded from the QWI. In addition, the QWI report job aggregates by workplace, so workers with multiple employers can be counted more than once.

[2] Quarterly earnings divided by three.

374. We apply the same exclusion to both earnings populations and all MSA-level analyses. While excluding these MSAs meaningfully limits the influence of noisy earnings averages from small hiring groups, our baseline results remain similar when these MSAs are included.

To relate these earnings gaps to local labor-market conditions, we construct measures of tightness from local unemployment and job postings. Monthly, not-seasonally-adjusted MSA unemployment rates come from the Bureau of Labor Statistics' Local Area Unemployment Statistics (LAUS) (US Bureau of Labor Statistics 2018). While LAUS measures workers by residence, QWI assign jobs by workplace. This is unlikely to represent a problem, since the geographic overlap between the two populations is substantial: in the 2024 ACS, 93 percent of workers living in an MSA work in that same MSA (US Census Bureau 2024).[3]

Lightcast supplies online job postings identifying the date, location, occupation, industry, and recruiting employer (Lightcast 2026).[4] We count newly posted advertisements and use these counts as a proxy for vacancies. Unlike new postings, vacancies measure the stock of positions open at a given point in time. Despite this distinction, Lightcast postings and job openings reported in the Job Openings and Labor Turnover Survey (JOLTS) (US Bureau of Labor Statistics 2026b) closely align across states: the correlations exceed 0.97 both in log levels and after removing common annual movements and 0.87 after removing persistent state heterogeneity (Appendix Figure C.2). This agreement supports the use of postings as a measure of local recruiting activity.

We use Lightcast's detailed information in three additional ways. First, we construct measures of minority-relative occupational tightness to assess whether the relationship between tightness and earnings gaps reflects stronger demand for jobs in which minorities are concentrated. Second, we use full posting texts to identify diversity, equity, and inclusion (DEI) content and recruitment for specialized DEI roles to examine the role of inclusion-oriented recruiting in the relationship between earnings gaps and tightness. Third, we use information on employers to construct their shares of recruiting within detailed occupations and measure local recruiting concentration, following Azar et al. (2020) and Choi and Marinescu (2024). We use this measure to study whether monopsony power affects the earnings-gap response to tightness. We provide further details on the construction of these measures in the relevant sections.

[3]The calculation uses national estimates from the 2024 one-year ACS, Table B08016, for workers ages 16 and older living in metropolitan areas, including those working from home.

[4]We use the Census Bureau's 2023 metropolitan delineations to assign job postings to MSAs using their reported counties.

We use CPS and ACS public-use microdata (Flood et al. 2025; US Census Bureau 2026c) to adjust for compositional change, since the QWI do not report all the individual characteristics needed for these adjustments. Both report education, age, sex, marital status, and employment, allowing us to assess whether local workforces change as markets tighten. The ACS's larger local samples and information on immigration history support more detailed composition adjustments than are feasible with our CPS extract. The CPS, in turn, provides monthly labor-market information and longer historical coverage.[5] Coverage and frequency differ across sources; we specify the sample and timing used in each exercise.

We use the Census Bureau's Job-to-Job Flows (J2J) and origin–destination statistics (US Census Bureau 2026b) to distinguish two channels that could explain the narrowing of minority earnings gaps: improved access to employment or movement to better-paying jobs. These quarterly aggregates report direct employer changes and transitions into and out of nonemployment by race and ethnicity, industry, and location. Unlike hiring counts alone, they distinguish entry into employment from movement between employers. The origin–destination tables identify the industries and MSAs of the jobs being left and entered and report earnings before and after the move for the same mover cohorts. This allows us to ask whether minorities reach better-paying destinations and obtain larger earnings gains, rather than simply making more moves.

Finally, we supplement these data with state minimum-wage histories (Vaghul and Zipperer 2022), updated using Economic Policy Institute and official state sources. These measures, deflated using the BLS CPI-U (US Bureau of Labor Statistics 2026a), allow us to assess whether minimum-wage changes account for the relationship between tightness and minority earnings gaps.[6]

4 Empirical Framework

We estimate the effect of local labor-market tightness on the earnings gap between white and minority workers, for metropolitan area m in quarter t:

$$G_{mt} \;=\; \beta\, T_{mt} \;+\; \sigma_{m,q(t)} \;+\; \delta_t \;+\; \gamma_m t \;+\; \varepsilon_{mt}, \tag{1}$$

[5] For the CPS-based analyses, we use the Census Bureau's 2003, 2013, and 2023 metropolitan delineation files, together with IPUMS CPS documentation on changes in metropolitan identifiers, to match survey areas to the geographic definitions used in our local tightness data.

[6] Appendix F.3.1 details their construction. Additional data used in supplementary exercises are described in the corresponding appendices.

where $G_{mt} = 100\left(\ln \overline{w}_{mt}^{W} - \ln \overline{w}_{mt}^{M}\right)$ is the minority earnings gap defined in Section 3. We use the log difference multiplied by 100 to approximate the percentage earnings gap, so changes in the gap are interpreted in percentage points. Here, T_{mt} is a measure of local labor-market tightness, and $\sigma_{m,q(t)}$, δ_t and $\gamma_m t$ are MSA-by-quarter-of-year fixed effects, calendar-quarter fixed effects and local linear trends, with $q(t)$ the quarter of the year in which t falls. Negative values of β indicate a narrowing of the earnings gap in tighter markets. Standard errors are clustered by metropolitan area.

The MSA-by-quarter-of-year fixed effects, $\sigma_{m,q(t)}$, absorb persistent differences across markets and market-specific seasonality. The MSA component removes cross-market variation in factors such as industry and occupational composition that may jointly shape average labor-market tightness and the minority earnings gap. The seasonal component accounts for recurring local fluctuations in labor demand—for example, harvest-related demand in agricultural areas—whose timing and magnitude differ across markets and may induce comovement between tightness and the earnings gap. Identification therefore comes from deviations in tightness from the level typically observed in the same market and quarter of the year. Calendar-quarter fixed effects remove the national business cycle and any other conditions common to all markets in a quarter.

Our preferred specification, equation (1), combines MSA-by-quarter-of-year and calendar-quarter fixed effects with MSA-specific linear trends. We include these trends to absorb gradual local changes that may jointly affect labor-market tightness and minority earnings gaps. Metropolitan areas follow divergent structural trajectories (Diamond 2016; Ganong and Shoag 2017), while aging, rising educational attainment, immigration, and minority population growth can change workforce composition at different rates across markets. Allowing each metropolitan area its own trend accounts for these different trajectories.

Our data show that gradual changes in composition matter for the estimated response to tightness and that local trends largely absorb their influence. We assess this by comparing how the tightness estimate changes when we add local trends or direct composition controls. If trends absorb the relevant compositional variation, adding the controls to a specification that already includes trends should change the estimate little. To measure composition, we use four ACS measures of the local working-age population: mean age, the share holding a bachelor's degree or more, the foreign-born share, and the minority share.[7] Because these measures are available for 237 of the

[7]Each measure is person-weighted over civilian household residents aged 16–64 and enters with a one-year lag. We

374 metropolitan areas from 2012 onward, we compare estimates within this subsample. Panel A of Table D.1 starts from a specification with MSA-by-quarter-of-year and calendar-quarter fixed effects and adds the composition controls, local trends, or both, holding observations fixed within each earnings population.

The results support using local trends to account for the compositional changes relevant to the tightness estimate. Relative to the specification with fixed effects alone, adding composition controls reduces the estimated response by approximately 15 percent.[8] Adding trends produces similar reductions, while adding composition controls to the specification with trends changes the estimates by only around 1 to 2 percent.[9] Although these four measures do not exhaust local compositional change, their limited additional contribution suggests that trends absorb the measured demographic variation relevant to the estimated earnings-gap response. Including local trends nonetheless comes at a cost: they may absorb genuinely gradual earnings responses to tightening alongside potentially confounding changes (Meer and West 2016; Neumark et al. 2014).

Together, the fixed effects and local trends substantially narrow the variation used to estimate the tightness response. The full specification absorbs over 90 percent of the variation in earnings gaps for beginning-of-quarter jobs and over 80 percent for new stable hires, as well as more than 90 percent of the variation in tightness. Persistent differences across metropolitan areas account for most of the earnings-gap variation, while common quarterly movements at the national level are particularly important for tightness. Identification therefore rests on the remaining changes within each market, net of its recurring seasonal pattern, common quarterly conditions, and local linear trend.

Because the effects of tightness may vary across metropolitan areas, equal and employment weighting may produce different estimates. We give each observed MSA-quarter equal weight because our interest is in how minority earnings gaps respond to tightening across local labor markets, rather than in how tightness affects the average minority worker. Employment weighting instead

construct the measures by assigning residential PUMAs to the metropolitan area containing most of their residents under the 2023 metropolitan delineations. The Census Bureau released no standard ACS one-year file for 2020, so quarters in 2021 use the 2019 values.

[8] We express all changes as percentages of the same starting fixed-effects-only estimate, rather than of the estimate immediately before each adjustment, so that the adjustments can be compared on a common scale.

[9] An alternative is to include state-by-year fixed effects, building on the literature's use of regional time controls to account for geographically uneven economic conditions (Allegretto et al. 2011; Dube et al. 2010). However, state-by-year fixed effects are considerably less effective at absorbing the relevant compositional changes than trends: adding composition controls to a specification already including state-by-year fixed effects changes the estimates by around 11 percent (Panel B of Table D.1).

gives greater weight to markets with more minority jobs, even though those jobs face the same measured tightness within each market-quarter. With all MSAs included, employment weighting assigns more than half of the total weight to the 20 largest markets, but only about 6 percent to the 200 smallest.

5 Core Results

5.1 Baseline Estimates

Tighter local labor markets narrow the minority earnings gap, with larger reductions among new stable hires. Table I reports the estimated effect of local labor-market tightness on the minority earnings gap, for jobs held at the beginning of the quarter and for new stable hires. Tightness is measured by minus the LAUS unemployment rate (columns 1 to 3) or by the log ratio of job postings to unemployment (columns 4 to 6).

In our main specification using unemployment (column 3), a 1-percentage-point decline in local unemployment narrows the minority earnings gap by approximately 0.2 percentage points for beginning-of-quarter jobs and 0.3 percentage points for new stable hires.[10] The same specification with the vacancy-to-unemployment ratio as our measure of tightness (column 6) yields similar results: doubling this ratio narrows the earnings gap by 0.3 and 0.4 percentage points, respectively.[11]

The finding that tighter local labor markets narrow the minority earnings gap also holds in less restrictive specifications, although the estimated magnitude varies. Because our preferred controls may absorb genuine earnings responses alongside confounding variation, we compare specifications to assess how these adjustments affect the estimated response. The alternative specifications retain MSA-by-quarter-of-year fixed effects alone (columns 1 and 4) or add state-by-calendar-quarter fixed effects (columns 2 and 5). The estimated responses remain negative and statistically significant across most specifications, with generally larger magnitudes than under our preferred specification. This pattern is consistent with compositional changes contributing, in less saturated specifications, to the estimated narrowing of minority earnings gaps as markets tighten (Section 4).

Finally, tightening narrows minority earnings gaps more strongly among new hires than among

[10] For beginning-of-quarter jobs, the effect amounts to a reduction of approximately 0.5 percent relative to the estimation-sample mean earnings gap of 40.3 percentage points. For new stable hires, the corresponding reduction is approximately 1.3 percent relative to a mean gap of 23.6 percentage points.

[11] These estimates give each MSA-quarter equal weight. Table E.1 reports the same specifications weighting by minority job counts and excluding the five largest metropolitan areas.

TABLE I: TIGHTNESS AND THE MINORITY EARNINGS GAP

	$-U$			$\ln(V/U)$		
Outcome: W - M earnings gap	MSA seasonal FE only	Add state-quarter FE	Add date FE and trends	MSA seasonal FE only	Add state-quarter FE	Add date FE and trends
Population: Beginning-of-quarter jobs						
Tightness	-0.61^{***}	-0.76^{***}	-0.21^{***}	-1.73^{***}	-0.57^{***}	-0.26^{**}
	(0.04)	(0.12)	(0.05)	(0.09)	(0.21)	(0.11)
Population: New stable hires						
Tightness	-1.17^{***}	-0.55^{***}	-0.30^{***}	-3.14^{***}	-0.34	-0.41^{**}
	(0.04)	(0.13)	(0.08)	(0.10)	(0.25)	(0.18)

Notes: Each cell reports a separate regression of the current-quarter minority earnings gap on the indicated tightness measure over 2010Q3–2025Q2. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline in the $-U$ columns and a doubling of postings relative to unemployment in the $\ln(V/U)$ columns. All specifications include MSA-by-quarter-of-year fixed effects. Columns 2 and 5 add state-by-calendar-quarter fixed effects; columns 3 and 6 instead add calendar-quarter fixed effects and MSA-specific linear trends. All estimates are unweighted and use 21,613–22,158 observations in 368–374 MSAs. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

beginning-of-quarter jobs. In our preferred specification, the response for new stable hires is about 40 percent more negative than for beginning-of-quarter jobs. This pattern is consistent with the substantial responsiveness of entry wages to labor-market conditions documented by Martins et al. (2012). New stable hires account for only about a tenth of minority jobs in a quarter, however, suggesting that their stronger response alone is unlikely to explain the earnings compression observed for the broader stock of jobs.

5.2 The Choice of Tightness Measure

The effect of tightness on the minority earnings gap in Table I is negative and significant under both the unemployment rate and the vacancy-to-unemployment ratio. Therefore, the finding that tightness reduces earnings gaps does not depend on which of the two measures we use.

Because unemployment enters both measures, their agreement in Table I is at least partly mechanical. We isolate what the vacancy side contributes by adding log postings per labor-force participant to equation (1) alongside the unemployment rate. Table E.3 places that specification beside the baseline. The unemployment coefficient is unchanged for both earnings outcomes, while neither recruiting coefficient is distinguishable from zero. Once the fixed effects, in particular the national movements, are absorbed, local variation in recruiting activity carries no information about the gap beyond the unemployment rate.

A null on the vacancy side could be an artifact of the job postings data we use, either because online job postings are a poor measure of vacancies at the local level or because their coverage grows non-linearly within MSA over the sample as advertising moves online, which would make postings rise without any change in vacancies.

We conduct two analyses that argue against both explanations. The first, in Table E.4, replaces online postings with the official vacancy count published in JOLTS, and repeats the estimation of Table E.3 at the state level, the most granular geography for which that count exists. The coefficient on the vacancy side is again indistinguishable from zero with either source of vacancies, although the unemployment side is estimated imprecisely at this geographic level. The second compares how fast the two sources grow. If expanding data coverage were driving the rise in postings, we would expect faster growth in online job postings than in official openings. However, that's not the case. Postings per new stable hire rise about 5.8 percent a year over the sample and official openings per new stable hire about 5.7 percent, with the ratio between the two showing no trend distinguishable from zero. While we find no national difference, we observe some differences across states and industries between the two measures, which prevents us from treating online job postings as a valid measure at the local level. Nevertheless, these analyses suggest that the null effect on the vacancy side of tightness is a feature of local labor markets rather than an artifact of our postings measure.[12]

This result does not contradict the literature on tightness and wages. The vacancy-to-unemployment ratio is the standard summary of tightness in work on aggregate wage growth, where it is formed from officially counted vacancies and set against a national average wage (Benigno and Eggertsson 2023). Domash and Summers (2022) enter the two sides separately and find that the vacancy side retains predictive power of its own over long samples. That evidence establishes an effect on wage growth, and is uninformative about how that effect differs between white and minority workers. A measure can be highly predictive of aggregate wage growth, yet uninformative about an earnings gap if it primarily captures wage pressure common to both groups: the gap changes only when white and minority earnings respond differently to that pressure. Thus, vacancies may raise earnings broadly without predicting convergence across groups unless their incidence or earnings effects differ between them. Moreover, their estimates rest on aggregate fluctuations over time, while ours set those aside and use variation across markets. Using the unemployment

[12] Equation (1) also addresses coverage growth by construction: within a market, coverage that grows at a constant rate enters the postings measure in logs as a linear time term, and the trends absorb it exactly.

rate to measure tightness is consistent with previous work on racial earnings gaps. For instance, Aaronson et al. (2019) use the distance between the national unemployment rate and the natural rate, while Autor et al. (2024) build a state index from unemployment and employer-to-employer separations.

Therefore, we use the unemployment rate as our measure of tightness in the remainder of the paper: the two measures are equally informative for the gap, and the unemployment rate is measured consistently across markets. Moreover, our choice is in line with the most closely related literature.

5.3 Nonlinearity in the Level of Slack

The constant slope in equation (1) imposes that a decline in unemployment from 9 to 8 percent carries the same effect as a decline from 4 to 3. There is reason to question it: at high unemployment, i.e., when labor-market slack is ample, an employer filling a vacancy can draw on a deep pool of available workers, whereas in a market already near full employment the same additional tightening may require raising pay or reaching further into the applicant queue (Benigno and Eggertsson 2023; Phillips 1958).

Several channels could make the response to tightening differ by group. Consider hiring standards: employers raise education and experience requirements when workers are plentiful and relax them when workers are scarce (Modestino et al. 2020). If the applicants those requirements screen out are disproportionately minority, relaxing them as markets tighten will affect the gap, and will affect it most where slack is already low. Employer wage-setting power and the difficulty of changing jobs may work the same way, since both operate through a worker's outside option. When labor is abundant, that option is weak for everyone, especially workers with few alternative employers. As labor becomes scarce, alternatives expand, with the greatest increase among those who initially had the fewest. National evidence shows this kind of state dependence: racial gaps in employment narrow more once unemployment falls below the natural rate (Aaronson et al. 2019). Here, we examine whether the minority earnings-gap response to tightening varies with local labor-market slack. We then explore the potential roles of job mobility and employer wage-setting power in Sections 7 and 9.

In Table II we relax the linearity assumption imposed in equation (1) by adding to our main specification a quadratic and a cubic term in tightness. The inclusion of these terms is motivated by

the original Phillips curve, which is convex in the level of unemployment (Phillips 1958), and by wage-bargaining models in which wages rise linearly with the vacancy-to-unemployment ratio (Pissarides 2000). The polynomial approximates that shape without imposing a particular direction of change in the slope. The fitted slope is evaluated at unemployment rates of 8.9, 5.6 and 2.9 percent, which correspond to the 10^{th} percentile, the mean, and the 90^{th} percentile of the tightness distribution.[13]

TABLE II: THE EFFECT OF TIGHTNESS AT COLD, AVERAGE, AND HOT MARKET CONDITIONS

Outcome: W - M earnings gap	Cold U = 8.87%	Difference: avg. – cold	Average U = 5.60%	Difference: hot – avg.	Hot U = 2.93%	Difference: hot – cold
Panel A. Tightness slopes at each market condition						
Population: Beginning-of-quarter jobs						
Tightness	-0.20^{***}	—	-0.28^{***}	—	-0.36^{***}	—
	(0.06)		(0.08)		(0.12)	
Population: New stable hires						
Tightness	-0.22^{**}	—	-0.40^{***}	—	-0.62^{***}	—
	(0.09)		(0.12)		(0.18)	
Panel B. Differences between the tightness slopes						
Population: Beginning-of-quarter jobs						
Tightness	—	-0.08	—	-0.09	—	-0.16
		(0.05)		(0.07)		(0.12)
Population: New stable hires						
Tightness	—	-0.18^{***}	—	-0.23^{***}	—	-0.41^{***}
		(0.07)		(0.08)		(0.15)

Notes: The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Panel A reports fitted responses to a 1-percentage-point unemployment decline from a cubic specification in tightness; Panel B reports differences between these responses. Entries are in percentage points: Panel A reports changes in the earnings gap and Panel B differences between those changes. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use 21,990–22,158 observations in 374 MSAs over 2010Q3–2025Q2. Cold and hot conditions correspond to the 90^{th} and 10^{th} percentiles of unemployment outside 2020; average conditions use the full-sample mean. Panel B subtracts the cooler-market response from the hotter-market response in the indicated order. Standard errors in parentheses use the MSA-clustered covariance matrix, including covariance between responses. Stars test a zero response in Panel A and a zero difference in Panel B: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The response steepens as slack is exhausted. For a 1-percentage-point unemployment decline, the estimated narrowing of the gap for jobs held at the beginning of the quarter rises from 0.2 percentage points in cold markets to 0.3 at average conditions and 0.4 in hot markets, and for new stable hires from 0.2 to 0.4 to 0.6 percentage points. Gains from a high-pressure economy are therefore larger where markets are already tight.

[13]The two percentiles are taken excluding 2020 while the mean is taken over the whole window, so that Table II and Table E.2, which excludes 2020, are evaluated at the same three conditions.

The effects on the two job populations diverge as slack disappears. Compared to the response for all jobs, new hires' response is the same in cold markets and almost doubles in hot ones. The stronger response of entry pay is therefore a feature of tight markets rather than a constant property of the relationship.

Panel B tests whether the slopes differ across conditions, using their estimated covariance. Moving from cold to hot markets increases the estimated earnings-gap compression by 0.2 percentage points for jobs held at the beginning of the quarter and by 0.4 percentage points for new stable hires, with the differences more precisely estimated for new hires. These differences are sizable relative to the average effect of tightening: the additional compression from cold to hot markets exceeds that effect for new stable hires and amounts to about three-quarters of it for beginning-of-quarter jobs.

5.4 Stability Across Historical Periods

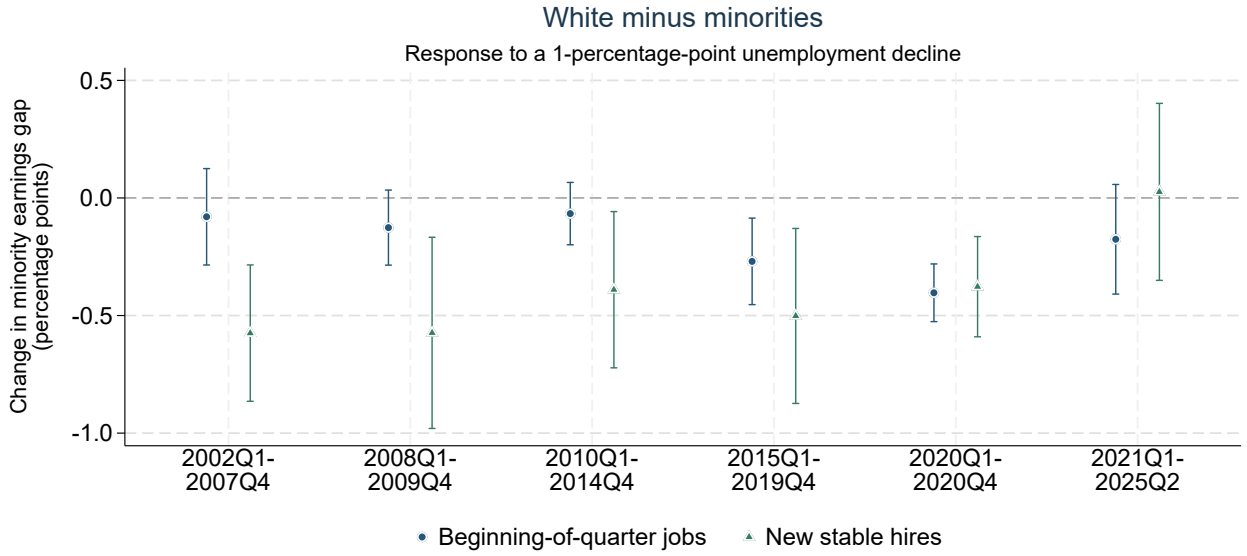
The overall relationship we estimate between tightness and earnings-gap compression could be driven by one or a few historical episodes in which unobserved local developments affected earnings gaps and were correlated with tightness. We find, however, that this relationship recurs across distinct periods, including both expansions and contractions, arguing against an explanation based on episode-specific confounding.

Figure I shows the relationship between tightness and minority earnings gaps across six different periods in our sample. We estimate period-specific tightness responses jointly for each earnings population, extending the sample back to 2002 to use as much history as the data allow. Each period corresponds to a distinct macroeconomic episode: the expansion of 2002–2007, the recession of 2008–2009, the slow recovery through 2014, the more robust expansion of 2015–2019, the pandemic year, and the expansion of 2021–2025. MSA intercepts and linear trends vary across windows, while MSA-by-quarter-of-year effects are common across periods. Cross-window magnitudes are not directly comparable: shorter windows retain less identifying variation after fitting local trends, while geographic coverage expands over time as states enter the QWI.

The estimate for jobs held at the beginning of the quarter is negative in every window, including the periods preceding our main estimation sample. Across these periods, a 1-percentage-point decline in unemployment narrows the minority earnings gap by approximately 0.1 to 0.4 percentage points. The estimates for new stable hires are generally more negative than those for all jobs

and are also broadly consistent across periods. Although the shorter estimation windows leave greater uncertainty around the individual estimates, taken together, these estimates indicate that the relationship between tightness and minority earnings gaps persists over time and is not driven by a single historical episode.

FIGURE I: THE EFFECT OF TIGHTNESS ON THE MINORITY EARNINGS GAP, BY PERIOD



Notes: Each point reports the period-specific response of the current-quarter minority earnings gap to tightness, from one regression per earnings population estimated jointly across periods. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects, with MSA effects and MSA-specific linear trends fitted within each period. Each period is a distinct macroeconomic episode. All estimates are unweighted and use 34,590 observations for beginning-of-quarter jobs and 34,138 for new stable hires in 374 MSAs over 2002Q1–2025Q2, with the all-period MSA exclusions described in Table D.1. Period-specific counts range from 1,488 to 8,739 and from 1,482 to 8,514, respectively, as metropolitan coverage expands when states enter the QWI. Vertical bars show 90 percent confidence intervals based on MSA-clustered standard errors.

6 Assessing the Role of Composition and Policy

The narrowing of minority earnings gaps in tighter labor markets could reflect changes in worker composition, a shift in vacancy composition toward occupations in which minorities are concentrated, or changes in concurrent public and employer policies during our sample period. We first assess whether changes in worker composition or selective geographic mobility explain this narrowing, then examine whether shifts in occupational demand account for the aggregate-tightness relationship. Finally, we examine whether changes in minimum wages or employers' DEI engagement explain the earnings compression associated with tightening.

6.1 Worker-side Composition

Our earnings gap outcome compares the mean earnings of white and minority workers employed in an MSA, so their earnings could converge if the characteristics of the two workforces changed differently as markets tightened. Accounting for observed workforce and population characteristics leaves the estimates unchanged. Similarly, selection into cross-MSA employer changes based on previous earnings does not contribute to the narrowing of the gap.

To assess whether changes in demographic composition explain the effects of tightness on earnings gaps, we add MSA-level controls constructed from the CPS and ACS to our specification, retaining QWI earnings gaps as the outcome. To capture earnings-related differences in the employed workforce, we first add white-minus-minority differences in age, sex, education, and marital status.[14] To account for the importance of migration during our period of study and its differential relevance across areas, the next ACS block adds immigration history and English proficiency to the previous basic demographics. The lack of comparable immigration information in our CPS extract precludes the same adjustment there. Finally, we add all-race population shares describing education, age, and, in the ACS, foreign birth. This last set of controls captures broader changes in local labor supply and MSA structure that could affect both tightness and earnings gaps. Because of the limited geographic detail in the public-use CPS and ACS files, these comparisons are restricted to fewer MSAs than the main analysis. Specifically, geographic identification and matching requirements restrict the exercise to 260 MSAs in the CPS and 237 MSAs over 2012–2024 (excluding 2020) in the ACS. For each survey, a no-controls baseline estimate on the same restricted sample provides the benchmark against which we compare the adjusted estimates, so that sample differences are not attributed to the controls. Appendix F.1.1 provides additional details on the covariates, survey populations, and timing.

The baseline estimate barely changes across the adjustments, as shown in Table III. In both the CPS- and ACS-restricted samples, a 1-percentage-point unemployment decline still narrows the minority earnings gap by about 0.25 percentage points for beginning-of-quarter jobs and 0.4 percentage points for new stable hires, with the estimates essentially unchanged after adding controls. Observed changes in workforce and population composition therefore do not account for the narrowing of the earnings gap with tightness.

Demographic controls account for observed workforce differences, but workers moving between

[14] The ACS controls also include differences in disability status.

TABLE III: MINORITY EARNINGS GAPS AND WORKER COMPOSITION

Outcome: W - M earnings gap	Baseline	Add demographic differences	Also add immigration differences	Also add aggregate population shares
Panel A. ACS workforce and population controls				
Population: Beginning-of-quarter jobs				
Tightness	-0.26^{***}	-0.26^{***}	-0.24^{***}	-0.25^{***}
	(0.09)	(0.08)	(0.08)	(0.08)
Population: New stable hires				
Tightness	-0.42^{***}	-0.44^{***}	-0.42^{***}	-0.43^{***}
	(0.14)	(0.14)	(0.14)	(0.14)
Panel B. CPS workforce and population controls				
Population: Beginning-of-quarter jobs				
Tightness	-0.24^{***}	-0.25^{***}	—	-0.25^{***}
	(0.06)	(0.06)		(0.06)
Population: New stable hires				
Tightness	-0.36^{***}	-0.36^{***}	—	-0.36^{***}
	(0.10)	(0.10)		(0.10)

Notes: Each cell reports a separate regression of the current-quarter minority earnings gap on tightness. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted. Controls, constructed using survey weights, accumulate within each panel; columns share a common sample within each earnings population. Workforce and immigration controls are white-minus-minority differences; population controls are all-race shares. The immigration adjustment is available only in the ACS. Panel A covers 2012–2019 and 2021–2024, with 11,204–11,226 observations in 237 MSAs; annual controls are assigned to each year's QWI quarters. Panel B covers 2010Q3–2025Q1, with 12,068–12,099 observations in 260 MSAs. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

MSAs may also differ in unobserved characteristics that affect their earnings. An inflow of relatively high-earning minority workers or an outflow of relatively low-earning minority workers could narrow the local gap without improving earnings for other workers. To assess the role of this type of selection, we focus on movers' previous earnings, which capture differences reflected in pay that the demographic controls need not capture. We focus on earnings before the move because earnings at the new job may already reflect improved opportunities: attributing those gains to composition would credit the arrival of new workers with an improvement that the move itself produced.[15] We use J2J origin–destination data, which report mean earnings before and after cross-MSA employer changes among continuously employed workers. Panel A of Table IV shows that employer-to-employer moves across MSAs, especially incoming moves, increase more for minorities than for whites as markets tighten. This suggests that selective cross-MSA sorting

[15]Previous earnings also reflect the origin employer and local labor-market conditions, so they summarize movers' prior earnings positions rather than provide a pure measure of worker productivity.

could be particularly relevant to earnings-gap compression.

Panel B compares movers' previous earnings with the mean earnings of their own group in the relevant local workforce: the destination for arrivals and the origin for departures. As markets tighten, the previous earnings of movers decline relative to these benchmarks. Selection into cross-MSA moves thus shifts toward lower-paid workers within each group, with no clear difference between whites and minorities. This similarity, however, does not by itself imply that the flows have the same effect on minority earnings gaps. Their contribution to each group's mean also depends on how many workers move and how far their earnings lie from that mean.

To quantify those contributions, we combine the number of movers with their previous earnings. For departures, we compare the pre-move local mean with the mean after removing departing workers and their earnings. For arrivals, we compare the mean earnings of the destination's other workers with the mean after adding arrivals at their previous pay. The arrival comparison therefore values workers at origin earnings but measures their contribution at the destination, without treating earnings gains at the new job as a change in composition. We calculate each contribution separately for whites and minorities, take the difference between the groups, and add the arrival and departure contributions to obtain their combined effect on the earnings gap. Appendix F.1.2 provides the construction details and timing conventions.

The accounting shows that arrivals and departures affect the two groups' mean earnings differently, but their contributions to the gap cancel out (Panel C). For a 1-percentage-point unemployment decline, the arrival contribution, valued at previous pay, widens the minority earnings gap by about 0.005 percentage points. Departures work in the opposite direction: removing leavers narrows the gap by roughly the same amount. The combined response of the minority earnings gap is approximately -0.0001 percentage points, with a standard error of 0.0030 percentage points. Thus, although cross-MSA job mobility increases more for minorities and the earnings position of movers changes with tightness, the net contribution of these flows to the earnings gap is small and statistically insignificant. Together with the demographic adjustments, this provides little support for changes in worker composition as an explanation of our earnings-gap finding.

6.2 Vacancy-side Composition

Changes in aggregate tightness could simply reflect a shift in the composition of vacancies toward occupations in which minorities are concentrated. Earnings gaps would then narrow because

TABLE IV: CROSS-MSA JOB MOBILITY AND ITS CONTRIBUTION TO LOCAL EARNINGS GAPS

	White	minorities	White minus minority
Panel A. Cross-MSA employer flows			
Outcome: Incoming cross-MSA moves / local stable jobs			
Tightness	$+0.02^{***}$	$+0.03^{***}$	-0.01^{***}
	(0.00)	(0.00)	(0.00)
Outcome: Outgoing cross-MSA moves / local stable jobs			
Tightness	$+0.01^{***}$	$+0.02^{***}$	-0.00
	(0.00)	(0.00)	(0.00)
Panel B. Movers' earnings relative to their group's local workforce			
Outcome: Pre-move earnings / destination pre-move mean			
Population: Arrivals			
Tightness	-0.37^{***}	-0.39^{***}	$+0.02$
	(0.14)	(0.15)	(0.16)
Outcome: Pre-move earnings / origin pre-move mean			
Population: Departures			
Tightness	-0.54^{***}	-0.51^{***}	-0.02
	(0.13)	(0.14)	(0.12)
Panel C. Flow contributions at pre-move earnings			
Outcome: Removing departing workers			
Tightness	$+0.006^{***}$	$+0.011^{***}$	-0.005^{**}
	(0.001)	(0.002)	(0.002)
Outcome: Arrivals valued at pre-move earnings			
Tightness	-0.007^{***}	-0.011^{***}	$+0.005^{**}$
	(0.002)	(0.002)	(0.002)
Outcome: Combined flow contribution at pre-move pay			
Tightness	-0.001	-0.001	-0.000
	(0.002)	(0.003)	(0.003)

Notes: Each cell reports a separate regression of the indicated outcome on tightness. Entries are responses to a 1-percentage-point unemployment decline; the final column directly estimates the white-minus-minority response. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use a common sample of 21,752 observations in 374 MSAs over 2010Q3–2025Q1. Panel A reports percentage-point changes in J2JOD employer moves between distinct MSAs, divided by same-group local stable jobs: destination jobs in $t+1$ for arrivals and origin jobs in $t-1$ for departures. These are workplace changes, not necessarily residential moves. In Panel B, the first two columns report approximate percent changes in movers' pre-move earnings relative to their own group's local pre-move mean; the final column reports the difference between these responses in percentage points. In Panel C, the first two columns report approximate percentage contributions to each group's local mean earnings from removing departures, adding arrivals at pre-move pay, and combining both flows; the final column reports the corresponding percentage-point contributions to the minority earnings gap. Black and Hispanic counts and earnings totals are pooled before constructing minority outcomes. Appendix F.1.2 provides the formulas and sample restrictions. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

demand rises disproportionately for these occupations, rather than because the local labor market tightens more broadly. Both explanations imply genuine relative earnings gains for minorities, but they point to different mechanisms through which tighter labor markets compress minority earnings gaps. We find that controlling for minority-relative occupational conditions leaves the estimated effect of aggregate tightness essentially unchanged. These findings suggest that broader tightening improves minorities' relative earnings through channels beyond stronger demand for minority-heavy occupations.

We construct minority-relative occupational tightness, R_{mt}, from recruiting and unemployment conditions in eight broad occupational groups, as follows:

$$\theta^r_{mt} = \frac{\sum_o s^r_{moy(t)} v_{mot}}{\sum_o s^r_{moy(t)} \widehat{u}_{mot}}, \qquad R_{mt} = \log \theta^M_{mt} - \log \theta^W_{mt}, \tag{2}$$

where r indexes the two groups, with M and W denoting minorities and whites. Within each MSA and occupation, recruiting intensity, v_{mot}, is Lightcast postings divided by employment in a fixed baseline period, making recruiting comparable across occupations of different sizes. We approximate local occupational unemployment, $\widehat{u}_{mot}$, by multiplying the MSA unemployment rate from LAUS by the ratio of the occupation's national unemployment rate to the national average, with both national rates calculated from the CPS. This preserves local differences in overall unemployment while allowing conditions to differ across occupations. We average recruiting intensity and unemployment separately using each group's time-varying ACS employment shares by occupation, $s^r_{moy(t)}$, which gives greater weight to occupations where more of its members work. Each group's tightness index, θ^r_{mt}, is the ratio of these two averages, and minority-relative tightness is the difference between the groups' log indices.[16] A higher R_{mt} indicates more favorable occupational conditions for minorities relative to whites. The comparison between minorities and whites reflects the occupations in which each group works, not vacancies or local unemployment observed separately by group. Appendix F.2 details the underlying measures, occupational mappings, and updating rules.

We include this relative index alongside aggregate tightness in the main specification of Section 4. The comparison examines whether the earnings gap still narrows with broader local tightening

[16] An unemployment-only measure of relative occupational tightness is highly correlated with aggregate $-U$ (0.61 after removing MSA-by-quarter-of-year and calendar-quarter effects and MSA trends). We therefore use the relative occupational $\ln(V/U)$ measure, whose correlation with aggregate $-U$ is much lower (-0.04).

when relative occupational conditions are held constant, and whether improvements in those conditions have themselves an additional relationship with the gap. The limited geographic detail in the ACS restricts MSA coverage, and we also require a minimum number (100) of minority employment observations to construct the occupational shares. These requirements restrict this exercise to 212 MSAs over 2010Q3–2025Q2.[17] We therefore compare the joint estimates with the aggregate-only estimates on the same sample, rather than with the larger-sample headline coefficient.

Accounting for minority-relative occupational conditions leaves the estimated effect of aggregate tightness essentially unchanged (Table V). More favorable minority-relative occupational conditions are weakly associated with additional earnings compression among new stable hires, but the corresponding estimate for beginning-of-quarter jobs is not statistically significant. After including the index, a 1-percentage-point unemployment decline narrows the minority earnings gap by approximately 0.3 percentage points for beginning-of-quarter jobs and 0.4 percentage points for new stable hires, essentially unchanged from the estimates without the occupational index. Panel B makes the magnitudes of the estimates for both tightness measures comparable by scaling each coefficient by the corresponding measure's standard deviation after accounting for the other measure, fixed effects, and trends.[18] A one-standard-deviation improvement in minority-relative occupational tightness is associated with a narrowing of the minority earnings gap of approximately 0.1 percentage points for both beginning-of-quarter jobs and new stable hires. These responses are approximately 30 and 45 percent as large as the corresponding responses to aggregate tightness, respectively. Relative occupational demand therefore is insufficient to explain why broader tightening narrows the gap.

6.3 Policy Levers

Beyond changes in worker and vacancy composition, the earnings compression associated with tightening could reflect policies affecting pay and access to jobs. Higher minimum wages or greater employer commitment to inclusion could improve minorities' relative earnings and, if

[17] Geographic matching supports 240 MSAs. Requiring at least 100 unweighted minority employment records across the 2012–2014 ACS releases, to construct the occupational shares, leaves 212. The national CPS occupational rates do not further restrict MSA coverage.

[18] We regress each measure on the other tightness measure, fixed effects, and trends and use the residual standard deviation on the estimation sample. For the preferred pairing, these are approximately 0.718 percentage points for aggregate $-U$ and 0.054 log points for relative occupational tightness. Coefficients and standard errors receive the same scaling, leaving significance unchanged. These conditional responses are not additive contributions to overall compression.

TABLE V: AGGREGATE AND MINORITY-RELATIVE OCCUPATIONAL TIGHTNESS

Outcome: W - M earnings gap	Aggregate tightness: $-U$	Aggregate tightness: $\ln(V/U)$
Panel A. Baseline scaling		
Population: Beginning-of-quarter jobs		
Aggregate tightness	-0.27^{***}	-0.45^{***}
	(0.06)	(0.15)
Minority-relative occupational tightness	-0.79	-0.59
	(0.57)	(0.56)
Population: New stable hires		
Aggregate tightness	-0.39^{***}	-0.61^{***}
	(0.10)	(0.23)
Minority-relative occupational tightness	-1.62^{*}	-1.35
	(0.84)	(0.82)
Panel B. Standardized effects		
Population: Beginning-of-quarter jobs		
Aggregate tightness	-0.20^{***}	-0.14^{***}
	(0.04)	(0.05)
Minority-relative occupational tightness	-0.06	-0.05
	(0.04)	(0.04)
Population: New stable hires		
Aggregate tightness	-0.28^{***}	-0.19^{***}
	(0.07)	(0.07)
Minority-relative occupational tightness	-0.13^{*}	-0.10
	(0.07)	(0.06)

Notes: Each column within an earnings population reports a regression including aggregate and minority-relative occupational tightness. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use 12,548–12,572 observations in 212 MSAs over 2010Q3–2025Q2; both columns use identical observations within each earnings population. Panel A scales aggregate $-U$ to a 1-percentage-point unemployment decline and log-V/U measures to a doubling. For minority-relative tightness, this is a doubling of the minority-to-white occupational-tightness ratio. Panel B scales each coefficient by its regressor's standard deviation after accounting for the other tightness measure, fixed effects, and trends. The occupational indices combine recruiting and unemployment conditions across eight occupational groups using annually updated employment shares separately for white and minority workers; Appendix F.2 details their construction. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

they coincide with tighter markets, partly explain this relationship. We examine these two possibilities by adding minimum wages and posting-based measures of DEI engagement to our main specification. Accounting for either leaves the estimated effect of tightness unchanged.

6.3.1 Minimum Wages

Minimum-wage increases can narrow minority earnings gaps because minorities are disproportionately represented among workers whose pay is affected (Wursten and Reich 2023). State and

local wage floors rose substantially during the 2010s (Aaronson et al. 2020), and previous evidence also links higher minimum wages to stronger labor-market conditions (Başkaya and Rubinstein 2012).[19] If these policy changes coincide with tightening, part of the earnings compression could therefore reflect higher wage floors rather than improved labor-market conditions.

To assess this possibility, we add the real statutory minimum wage to the preferred specification in Section 4. The measure uses the higher of the state and federal minimum wages, averaged within each quarter and expressed in 2024 dollars. For MSAs crossing state borders, we combine state values using fixed employment shares, so that the measure does not change simply because employment shifts toward the part of the MSA with a higher wage floor. We compare the baseline and adjusted estimates on identical observations within each earnings population, retaining beginning-of-quarter jobs and new stable hires over 2010Q3–2025Q2. Appendix F.3.1 provides the construction details.

TABLE VI: EARNINGS-GAP COMPRESSION AFTER ACCOUNTING FOR MINIMUM WAGES

Outcome: W - M earnings gap	Baseline	+ Minimum wage
Population: Beginning-of-quarter jobs		
Tightness ($-U$)	-0.21^{***}	-0.22^{***}
	(0.05)	(0.05)
Real effective minimum wage	—	-0.18
		(0.11)
Population: New stable hires		
Tightness ($-U$)	-0.30^{***}	-0.31^{***}
	(0.08)	(0.08)
Real effective minimum wage	—	-0.34^{**}
		(0.14)

Notes: Each column within an earnings population reports a regression of the current-quarter minority earnings gap on tightness, with or without the real minimum wage. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline or a one-dollar increase in the minimum wage, in 2024 dollars. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use 21,990–22,158 observations in 374 MSAs over 2010Q3–2025Q2. Both columns use identical observations within each earnings population. The wage floor averages applicable state/federal rates within each quarter; fixed employment shares combine state rates in cross-state MSAs. Local ordinances are not comprehensively captured. Appendix F.3.1 provides the construction details. The adjusted column's two coefficients come from the same regression. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Accounting for minimum wages leaves the earnings-gap response to tightness essentially unchanged (Table VI). For a 1-percentage-point unemployment decline, the estimated narrowing of the minority earnings gap remains 0.2 percentage points among beginning-of-quarter jobs and

[19] That analysis uses state-year observations for 1978–2007, outside of our period of analysis.

around 0.3 percentage points among new stable hires. Higher minimum wages are themselves associated with narrower gaps, particularly among new hires: a one-dollar increase in the real wage floor is associated with a 0.3-percentage-point narrower minority earnings gap among new hires. The corresponding estimate for beginning-of-quarter jobs is a statistically insignificant 0.2 percentage points. Minimum wages may therefore contribute to narrower earnings gaps, but their observed variation does not account for the earnings compression associated with tightening.

6.3.2 Diversity, Equity, and Inclusion

Employers' efforts to broaden recruitment and support underrepresented workers could also contribute to minority earnings compression. These initiatives became more prominent during our sample period, particularly following 2020 (Abdelwahed 2025; Indeed and Glassdoor 2022), and their scale appears sensitive to firms' economic circumstances (Ayas et al. 2023). Because stronger local demand improves those circumstances while also tightening labor markets, employers may have greater scope to expand these initiatives precisely when markets are tight. Therefore, part of the estimated tightening response could reflect changes in employers' commitment to inclusion. We assess this possibility by exploiting the full text of job postings to construct two novel measures of diversity, equity, and inclusion (DEI) engagement. Accounting for either leaves the earnings-gap response to tightening essentially unchanged.

To identify language describing inclusion commitments, practices, and specialized DEI roles we apply a rule-based algorithm to job-advertisement titles and bodies. After standardizing text formatting, the algorithm searches for predefined phrases and combinations of words. Some rules require words to appear close together—for example, language encouraging applications alongside a reference to underrepresented groups—rather than relying on isolated mentions of diversity or inclusion. The rules accommodate common wording variants and record whether qualifying language is present, rather than counting how often DEI-related words appear. Appendix F.3.2.1 describes the algorithm in detail.

We use the resulting classifications to construct two indicators of DEI engagement. Our first measure, substantive DEI content, captures employers' stated commitments or actions to promote inclusion. It flags language describing commitments to an inclusive workplace, outreach to underrepresented applicants, inclusive hiring practices, or support for employee resource groups. We do not count "boilerplate" legal language, including generic equal-opportunity and nondis-

crimination statements, as evidence of substantive DEI content. Our second measure, recruitment for specialized DEI roles, captures employers' decisions to advertise positions with a dedicated DEI function. It identifies explicit DEI functions or combinations of inclusion-related language and occupational terms in job titles, such as "diversity manager," while excluding specified unrelated uses, such as "financial equity." This measure captures advertisements for those roles, not all vacancies posted by employers that recruit for them. Across our posting sample, the algorithm identifies specialized DEI roles in 0.03 percent of advertisements and substantive DEI content in 6.6 percent of them. Both measures encompass inclusion beyond race and describe advertised engagement, not verified implementation or effectiveness.

Comparisons with human reviews lend support to both measures as indicators of advertised DEI engagement, with closer agreement on specialized roles than on substantive inclusion language. We compare the algorithmic classifications with human reviews of 600 postings, using a sample that oversamples algorithm-positive postings and edge cases. For specialized DEI roles, the algorithm agrees with human reviewers in 88 percent of cases. The corresponding agreement rate for substantive DEI content is 69 percent. Greater disagreement over substantive content is expected, given the more subjective nature of its classification. In the human-reviewed sample, the algorithm classifies more advertisements as substantive than human reviewers. The algorithm captures most human-identified substantive content but also identifies cases that reviewers do not accept. In particular, the algorithm counts explicit inclusion commitments even when no concrete initiative is described. Human reviewers appear to apply a stricter threshold, treating some of these commitments as generic employer language rather than substantive engagement. Appendix F.3.2.1 provides the review procedures and detailed comparisons.

External comparisons suggest that employers whose job postings more often contain substantive DEI content also show broader commitments to inclusion. We assess this relationship using commitments documented outside employers' advertisements. Among employers committed to pursuing Black Equity at Work certification—a program setting standards for workplace racial equity—substantive DEI content appears in approximately 35 percent of postings on average, compared with 6 percent among employers in similar industries and with comparable posting volumes. Substantive DEI content in job postings is also more prevalent among the career-services sponsors and partners listed by North Carolina A&T, a historically Black university.[20]

[20]Neither DEI measure clearly predicts smaller Black–white callback disparities in the independent correspondence experiment reported by Kline et al. (2024). However, the comparison includes only 92 employers, which may limit

Appendix F.3.2.2 reports these comparisons.

To assess whether local DEI engagement accounts for the earnings compression associated with tightening, we construct separate MSA-level shares of eligible advertisements containing substantive DEI content and advertising specialized DEI roles. We pool postings over the current and previous three quarters to smooth fluctuations in relatively uncommon DEI recruiting. We then add each measure separately to the preferred specification in Section 4. Because classified postings begin in 2010Q3, the four-quarter DEI shares restrict this comparison to 2011Q2–2025Q2. We re-estimate the baseline on the same observations so that changes from the baseline tightness estimate reflect the added control rather than sample differences. Appendix F.3.2.3 provides additional detail on the construction and scaling of the DEI measures.

TABLE VII: EARNINGS-GAP COMPRESSION AFTER ACCOUNTING FOR DEI ENGAGEMENT

Outcome: W - M earnings gap	Substantive DEI content	Specialized DEI roles
Population: Beginning-of-quarter jobs		
Tightness ($-U$)	-0.22^{***}	-0.22^{***}
	(0.05)	(0.05)
DEI control	-0.04	$+0.00$
	(0.05)	(0.05)
Population: New stable hires		
Tightness ($-U$)	-0.28^{***}	-0.28^{***}
	(0.09)	(0.09)
DEI control	$+0.09$	-0.03
	(0.09)	(0.10)

Notes: Each column within an earnings population reports a regression including tightness and the indicated DEI control; the two DEI measures enter separately. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline or a one-interquartile-range increase in DEI engagement. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use 20,536–20,694 observations in 368 MSAs over 2011Q2–2025Q2, with identical observations across columns within each earnings population. DEI shares pool eligible advertisements over the current and previous three quarters. Substantive content identifies inclusion commitments or actions beyond legal boilerplate; specialized roles identify recruitment for dedicated DEI functions. Reference interquartile ranges, from the 2017–2019 distribution across 374 MSAs, are 2.31 percentage points for substantive content and 2.12 advertisements per 10,000 for specialized roles. Appendix F.3.2.3 details construction and scaling. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Accounting for either DEI measure leaves the earnings-gap response to tightening essentially unchanged (Table VII). A 1-percentage-point unemployment decline is associated with a narrowing of the minority earnings gap of approximately 0.2 percentage points for beginning-of-quarter jobs and 0.3 percentage points for new stable hires. These estimates are the same as those obtained

our ability to detect such a relationship.

without DEI controls. The DEI measures themselves have little explanatory power and are insignificant. Therefore, changes in DEI engagement captured by our posting-based measures do not account for the narrowing of the minority earnings gap as labor markets tighten.

7 Transition Frequency and Earnings-Gap Compression

Workers change employers more often when labor markets tighten (Akerlof et al. 1988), and those moves on average reallocate them toward better-paying employers (Moscarini and Postel-Vinay 2018). Tightening may raise mobility unevenly across groups of workers, and a difference in those responses could account for the compression documented in Section 5. We assess whether minority transition rates rise more strongly than white transition rates as markets tighten, and then examine whether these relative changes can account for the effect of tightness on the earnings gap. Minorities' hiring and job-to-job transitions rise more strongly as markets tighten, but these increases do not explain why tightness narrows the earnings gap.

We measure transitions using LEHD's Job-to-Job Flows data. From these data, we construct separate rates of hiring from another employer or from nonemployment, and of separation to another employer or into nonemployment. Each J2J rate divides the quarterly flow by average main-job employment at the beginning and end of the quarter.[21] We also use the QWI to compare initial hiring with hiring that reaches stable status (i.e., employment relationships that continue through the following quarter). For each measure, the comparison is the white rate minus the minority rate, where the latter is constructed by weighting Black and Hispanic rates by each group's beginning-of-quarter employment in the QWI. Appendix G provides further details on the measurement of each rate and on the estimates behind each panel.

Panel A of Table VIII shows that tighter labor markets increase hiring more strongly for minority workers than for white workers. For hiring from another employer, a 1-percentage-point decline in local unemployment raises the hiring rate through direct employer changes by 0.10 percentage points for white workers and by 0.14 for minority workers. The contrast is sharper for the hiring rate from nonemployment, which rises only for minority workers. Both differential responses are precisely estimated.

The separation results reinforce this pattern. Separations to another employer rise by 0.11 per-

[21] Flows accrue over the whole quarter while employment is measured at a single point in time. Averaging the beginning and the end approximates average employment over the quarter, which is the conventional base for a quarterly flow rate.

TABLE VIII: TRANSITION RATES AND LOCAL LABOR-MARKET TIGHTNESS

Panel A. Job transitions and access to stable jobs	White	minorities	White minus minority
Outcome: Direct employer-change hires (EE)			
Tightness	$+0.10^{***}$	$+0.14^{***}$	-0.04^{***}
	(0.01)	(0.02)	(0.01)
Outcome: Direct employer-change separations (EE)			
Tightness	$+0.07^{***}$	$+0.11^{***}$	-0.04^{***}
	(0.01)	(0.01)	(0.01)
Outcome: Hires from nonemployment (NE)			
Tightness	$+0.03$	$+0.08^{***}$	-0.05^{***}
	(0.02)	(0.02)	(0.01)
Outcome: Separations into nonemployment (EN)			
Tightness	-0.04	-0.04	$+0.01$
	(0.03)	(0.03)	(0.01)
Outcome: Initial hires in entry quarter t			
Tightness	$+0.26^{***}$	$+0.48^{***}$	-0.23^{***}
	(0.02)	(0.06)	(0.04)
Outcome: Subsequent stable hires, recorded in t+1			
Tightness	$+0.09^{***}$	$+0.13^{***}$	-0.04^{***}
	(0.01)	(0.02)	(0.01)

Panel B. White-minus-minority transition responses by market temperature	Cold	Average	Hot
Outcome: Direct employer-change hires (EE)			
Tightness	-0.04^{***}	-0.10^{***}	-0.16^{***}
	(0.01)	(0.01)	(0.02)
Outcome: Direct employer-change separations (EE)			
Tightness	-0.04^{***}	-0.09^{***}	-0.13^{***}
	(0.01)	(0.01)	(0.02)
Outcome: Hires from nonemployment (NE)			
Tightness	-0.06^{***}	-0.04^{**}	-0.02
	(0.01)	(0.02)	(0.03)
Outcome: Separations into nonemployment (EN)			
Tightness	$+0.01$	-0.01	-0.04^{**}
	(0.01)	(0.01)	(0.02)

Notes: Each cell reports a response of the indicated transition rate to a 1-percentage-point unemployment decline, in percentage points. Panel A reports linear estimates for white workers, minorities, and their directly estimated difference. Panel B reports white-minus-minority responses from cubic specifications, evaluated at unemployment rates of 8.87, 5.47, and 2.93 percent; these benchmarks exclude 2020, but the regressions include it. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted. EE denotes direct employer changes, NE hires from nonemployment, and EN separations into nonemployment. J2J rates use average beginning- and end-of-quarter main-job employment as the denominator; minority rates use QWI employment weights. The QWI hiring outcomes align entry in quarter t: initial hiring is $100HirN_t/Emp_t$, and subsequent stable hiring is $100HirNS_{t+1}/Emp_t$. J2J estimates use 22,051 observations over 2010Q3–2025Q2; QWI estimates use 21,548 observations over 2010Q3–2025Q1. Both cover 374 MSAs. Appendix G provides definitions. Standard errors are MSA-clustered. Stars use t tests in Panel A and normal-reference tests in Panel B: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

centage points for minority workers, compared with 0.07 for white workers, with a precisely estimated difference. Separations into nonemployment, by contrast, decline similarly for both groups, though neither decline is statistically significant. Together, the separation results indicate that tightening increases moves between employers rather than exits from employment.

To assess whether the stronger minority hiring response extends to more durable employment relationships, we compare initial hiring with jobs that reach stable status in the following quarter. Both rates increase more for minority workers than for white workers, but the minority-white difference falls from 0.23 percentage points for initial hiring to just 0.04 for hiring that reaches stable status. Tightening therefore increases minority workers' entry into stable employment, but by substantially less than the initial hiring response would suggest.

Because tightening compresses earnings gaps more strongly in already-tight markets, Panel B examines whether differences in transition-rate responses between minorities and whites also depend on labor-market slack. The employer-change margins steepen as slack disappears: a 1-percentage-point decline in unemployment raises the hiring rate from another employer by 0.04 percentage points more for minorities than for whites in cold markets, compared with 0.16 percentage points more in hot markets. For hiring from nonemployment, minorities' additional response relative to whites falls from 0.06 percentage points in cold markets to 0.02 in hot ones. This may reflect minorities' greater likelihood of being unemployed, which leaves more scope for relative gains through hiring when slack is high, with this source of gains becoming less important as labor-market conditions improve. Thus, as slack disappears, minorities' transition rates respond more strongly to tightening relative to whites', driven by a growing employer-to-employer response that more than offsets a shrinking relative response in hiring from nonemployment.

Table IX examines whether the observed differences in transition rates explain why tightness narrows the minority earnings gap. To do so, we include the white-minus-minority differences in flow rates as controls in equation (1). For all jobs, the estimated narrowing is essentially unchanged under every set of flow controls. For new stable hires, the estimated narrowing declines from 0.29 percentage points in the baseline to 0.23 when controlling for the stable-hiring difference and to 0.21 when the employer-change difference is included as well.[22] Thus, minorities' greater gains in access to stable employment appear to partly explain why tightening narrows the earnings gap

[22]The flow rates are unavailable in a small number of MSA–quarters. For comparability across columns, we fix the sample in this exercise to the observations for which every rate is available. The restriction alters the new-hire baseline by 0.005 percentage points, from the 0.30 of Table I to 0.29.

TABLE IX: THE MINORITY EARNINGS GAP AFTER ACCOUNTING FOR TRANSITION RATES

	No flow controls	Controls for white-minus-minority flow differences			
Outcome: W - M earnings gap	Baseline	EE	EE, NE and EN	Stable hiring	EE + stable hiring
Population: Beginning-of-quarter jobs					
Tightness	-0.21^{***}	-0.25^{***}	-0.23^{***}	-0.21^{***}	-0.25^{***}
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Population: New stable hires					
Tightness	-0.29^{***}	-0.27^{***}	-0.26^{***}	-0.23^{***}	-0.21^{**}
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)

Notes: Each cell reports a separate regression of the current-quarter minority earnings gap on tightness and the indicated flow controls. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Entries are approximate percentage-point changes in the earnings gap for a 1-percentage-point unemployment decline. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and use 21,901–22,051 observations in 374 MSAs over 2010Q3–2025Q2. All columns use identical observations within each earnings population. Controls are white-minus-minority rate differences. EE is the direct employer-change separation rate, NE hiring from nonemployment, and EN separation into nonemployment. Stable hiring is measured in the current quarter, $100 HirNS_t / Emp_t$, unlike the forward-aligned outcome in Table VIII. Column 4 includes only stable hiring; column 5 combines it with EE. Appendix G.3 provides the construction details. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

among new hires.

Thus, minorities' stronger transition-rate responses do not explain why tightness narrows minority earnings gaps among all jobs, though they account for a small part of the response among new hires. Transition frequency, however, only captures how often workers move, not the types of jobs they move to, which we examine in Section 8.

8 Transition Quality and Earnings-Gap Compression

We next study whether tightness changes the quality of the jobs workers obtain. We first rank hiring destinations by pay across industry, firm size, and geography and compare how their ranks respond to tightening. We then account for where workers come from, asking whether an increase in tightness leads minorities to upgrade more than whites from lower-paying to higher-paying industries or metropolitan areas and whether such upgrading contributes to their relative earnings gains from changing employers. Overall, we find limited evidence of stronger upgrading for minorities across these broad pay categories. However, comparing the full hiring distributions across job types reveals a clear reduction in separation at hiring between whites and minorities, capturing a broader convergence in the jobs the two groups obtain than pay rankings alone reveal.

We focus first on hiring destinations, using the QWI's category earnings and initial-hiring counts.

We rank industries by pay to assess whether hiring shifts toward higher-paying sectors. To capture geographic differences within local labor markets, we also rank counties within each MSA to assess whether hiring shifts toward higher-paying counties. Industry, county, firm-size, and joint job-type rankings all use inflation-adjusted average earnings in stable jobs over the three preceding calendar years and are updated annually. For each measure, we calculate the average destination rank of initial hires. The comparison is the white mean rank minus the minority mean rank, with the latter weighting Black and Hispanic mean ranks by their respective hire counts.

Table X compares how destination pay ranks respond to tightening across groups. Panel A considers industry, firm size, and geography separately. The most substantial improvement occurs in firm size: a 1-percentage-point decline in unemployment raises the destination rank of initial hires by about 0.2 points for both groups, but the difference between their responses is not distinguishable from zero. Along the industry dimension, the corresponding increase is 0.05 points for white hires and 0.09 for minority hires, with a statistically significant difference of 0.05 points in favor of minorities. Neither group shows a clear shift toward higher-paying counties within the hiring MSA. Because a job destination can rank highly in one dimension but poorly in another, the first row of Panel B considers industry, firm size, and geography in a joint ranking of destinations, while the second row ranks industry-by-size job types within each MSA to capture shifts toward better-paying local jobs that differences in pay across locations alone may obscure. When job types are ranked across all MSAs, neither group's response is distinguishable from zero, whereas ranks rise similarly with tightness for both groups when job types are compared within the same MSA. Taken together, these measures provide very limited evidence that tightening leads minorities to upgrade more than whites across the broad job categories we observe.

To assess whether tightening reduces segmentation between minorities and whites at hiring, Panel C compares the full hiring distributions of the two groups across industry-by-size job types. This comparison can reveal convergence in the jobs the two groups obtain that their average pay ranks may conceal. We measure the distance between these distributions using the standard dissimilarity index: a value of zero indicates identical hiring distributions, while larger values imply greater separation. A 1-percentage-point decline in unemployment reduces the index by 0.13 points. Tightening therefore makes initial hiring for white and minority workers more similar in composition, with limited evidence of disproportionate movement of minority workers toward higher-paying categories. The convergence may instead occur through shifts between similarly

TABLE X: JOB DESTINATIONS AND SEPARATION BETWEEN MINORITIES AND WHITES IN HIRING

	white	minorities	white minus minority
Panel A. Individual dimensions of initial-hire destinations			
Outcome: Industry: initial-hire destination rank			
Tightness	$+0.05$	$+0.09^{**}$	-0.05^{***}
	(0.04)	(0.04)	(0.02)
Outcome: Firm size: initial-hire destination rank			
Tightness	$+0.21^{***}$	$+0.23^{***}$	-0.02
	(0.07)	(0.08)	(0.04)
Outcome: Geography: county destination rank within MSA			
Tightness	-0.02	-0.04	$+0.02$
	(0.16)	(0.17)	(0.03)
Panel B. Joint dimensions of initial-hire destinations			
Outcome: Industry, size, and geography jointly			
Tightness	-0.03	$+0.00$	-0.03
	(0.06)	(0.06)	(0.03)
Outcome: Industry and size jointly, within MSA			
Tightness	$+0.17^{***}$	$+0.22^{***}$	-0.04
	(0.06)	(0.07)	(0.03)
Panel C. Separation in initial hiring			
Outcome: Distance between white and minority hiring distributions			
Tightness	—	—	-0.13^{***}
			(0.04)

Notes: Each cell reports a separate regression of the indicated QWI initial-hiring outcome on tightness, measured in the receiving MSA. Entries are responses to a 1-percentage-point unemployment decline. Panels A and B use destination pay ranks on a zero-to-100 scale; their final column directly estimates the white-minus-minority response. Panel A considers industries, firm-size classes, and counties within the hiring MSA. All pay ranks use all-worker private-sector stable-job earnings over the three preceding calendar years, deflated using the CPI-U, weighted by stable-job counts, and updated annually. Each benchmark requires at least eight valid quarters and at least two in each reference year. County comparisons require at least two ranked counties within the MSA. Panel B ranks joint job types across MSAs or within each MSA. Mean ranks are weighted by initial-hiring counts, pooling Black and Hispanic counts for minorities; both groups must have at least 75 percent of initial hires in commonly available categories. Panel C measures separation between the groups as 100 times one-half the sum of absolute differences in white and minority hiring shares across industry–firm-size types; negative estimates indicate less separation. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and cover 2010Q3–2025Q2. Panel A uses 21,317 observations in 367 MSAs for industry, 21,648 in 374 for firm size, and 13,686 in 233 for counties. Panel B uses 11,859 observations in 300 MSAs; Panel C uses 12,591 in 317. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

ranked job types or offsetting movements across higher- and lower-ranked types.

Changes in hiring destinations as markets tighten, however, do not by themselves reveal the gains from changing jobs, since tightening may also change the jobs workers leave. A shift toward better-paying destinations need not imply greater upgrading if workers also move from better-paying origins. To account for both sides of a move, we use LEHD's Job-to-Job Origin-Destination data, which report flows and mean earnings before and after employer transitions. Unlike the initial-hiring comparisons above, this exercise follows workers between employers, including direct changes and hires following brief nonemployment.[23] Within each origin MSA, quarter, and group, we pool earnings using mover counts and then calculate the log change from mean origin to mean destination earnings. Table XI separates this observed gain into the amount associated with differences in pay between origin and destination industries or MSAs and a residual. We construct the industry- or MSA-pay component using the same calculation and mover weights, but replacing observed earnings with the corresponding categories' inflation-adjusted average earnings in stable QWI jobs over the three calendar years preceding the transition.

Table XI provides little evidence that tightening leads minority movers to upgrade more than white movers across industries or metropolitan areas. Once both origin and destination pay are taken into account, the industry-pay component points to slightly greater upgrading for minorities, consistent with the hiring-destination results, whereas the metropolitan-pay component is close to zero. Both responses are statistically insignificant. Accounting for where workers move from therefore provides little support for disproportionate minority upgrading across these categories as a source of earnings compression. The white-minus-minority difference in total earnings gains also shows no clear response to tightening. However, this result concerns differences between the groups and therefore does not imply that movers obtain no earnings gains from tightening.

Taken together, the destination and earnings-gain results provide limited evidence that tightness narrows minority earnings gaps through systematically greater minority upgrading across broad pay categories. Nevertheless, tightening makes the hiring distributions of the two groups more similar, which could reflect improvements in the jobs minorities obtain that help narrow earnings gaps without being captured by standard quality measures. These results leave open the possibility that tightening generates earnings improvements within broad job categories and in existing

[23] Both measures cover transitions between stable jobs, with earnings measured over a full quarter in the origin job before the move and in the destination job afterward.

TABLE XI: EARNINGS GAINS ACROSS EMPLOYER TRANSITIONS AND THE CONTRIBUTION OF BROAD PAY CATEGORIES

	Industry pay categories	Metropolitan pay categories
Population: Employer transitions		
Outcome: Observed W - M earnings-gain difference		
Tightness	−0.00	−0.00
	(0.07)	(0.07)
Outcome: Component accounted for by category pay		
Tightness	−0.05	−0.01
	(0.04)	(0.02)
Outcome: Remaining W - M earnings-gain difference		
Tightness	+0.05	+0.01
	(0.08)	(0.07)

Notes: Each cell reports a separate regression of the white-minus-minority difference in earnings gains on tightness. Entries are approximate percentage-point changes in the white-minus-minority earnings-gain difference for a 1-percentage-point unemployment decline. J2JOD mover cohorts include stable direct employer changes and stable hires following brief nonemployment. Within each origin MSA, quarter, and group, we sum mover counts and earnings totals before calculating mean origin and destination earnings; the observed gain is the difference between the logs of these means. Black and Hispanic counts and earnings totals are likewise combined before taking logs. The category-pay component applies the same calculation and mover weights to origin and destination category pay, using all-worker private-sector QWI stable-job earnings over the three calendar years preceding the transition. Earnings are deflated using the CPI-U and weighted by stable-job counts. Benchmarks are updated annually and require at least eight valid quarters and at least two in each reference year. Both endpoints of a move use the same transition-year reference. The residual is the observed gain minus this component. Both columns use identical origin–destination cells jointly defined by industry and MSA, retaining only cells with released, positive mover counts and earnings at both endpoints and available industry and metropolitan pay benchmarks. The same retained movers underlie all three outcomes in both columns. All six regressions use the same 21,682 eligible origin-MSA–quarters in 374 MSAs from the joint-cell construction. All specifications include origin-MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends. Estimates are unweighted and cover 2010Q3–2025Q1. MSA-clustered standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

jobs. Section 9 examines employer competition as a potential mechanism for such gains.

9 Monopsony Power

Competition among employers, not just overall labor demand, affects workers' opportunities to move to better-paying jobs and may shape the earnings gains minorities obtain relative to whites as markets tighten. We find that recruiting concentration limits the additional minority earnings-gap compression that tightening produces as labor markets heat up. In hot markets, in particular, the earnings-gap response to tightening is meaningfully weaker where recruiting is more concentrated.

Recruiting concentration may limit how strongly tightening increases the pressure on employers to improve pay. Tighter markets can improve workers' outside opportunities, making it harder to

recruit and retain workers at low pay, as discussed by Langella and Manning (2021). However, if hiring expands primarily among a few employers, stronger labor demand need not generate the same competition for workers as when recruiting is spread across many employers. Under these conditions, the same increase in tightness could produce smaller improvements in workers' ability to obtain better pay.

Concentration's implications for minority earnings gaps depend on which workers gain most from the competitive pressure created by tightening. In the job-ladder framework discussed by Autor et al. (2024), workers in lower-paying jobs may benefit particularly from the stronger employer competition associated with tightening because alternative offers are more likely to improve on their current pay, through moves to better-paying firms or raises by current employers. Because Black and Hispanic workers are disproportionately represented in lower-paying jobs (Ross and Bateman 2019), stronger competition could raise their earnings relative to whites.[24] Recruiting concentration could limit those relative gains by weakening the competitive pressure that tightening creates.

To proxy for local employer market power, we measure recruiting concentration among employers seeking workers in the same local occupation, following Azar et al. (2020). For each MSA and six-digit occupation, we calculate a quarterly Herfindahl–Hirschman index (HHI) from employers' shares of job postings. The index is higher when fewer employers account for a higher share of the recruiting. Our MSA-level measure is a posting-weighted average across occupations, using each occupation's average HHI over the preceding four quarters. This measure describes the recruiting environment both minorities and whites face, rather than concentration among the employers hiring from each group separately. Appendix I details the construction and benchmarks the measure against published concentration and pay patterns.

Because we aim to understand how market power shapes the earnings-gap response to tightening, we interact our measure of local employer market power with tightness.[25] We first estimate a linear relationship and then allow the role of concentration to vary with the degree of slack in the labor market, following equation (8). Our measure of concentration is constructed using postings from the preceding four quarters, so this comparison covers only the period 2011Q3–2025Q2. Table XII reports the linear estimates in Panel A. Panel B compares earnings-gap responses at low

[24] This channel can operate through differences across employers within occupations, consistent with our finding in Section 6.2 that minority-relative occupational conditions do not account for the aggregate-tightness relationship.

[25] Our specifications also include concentration as a separate control.

(p25) and high (p75) concentration in cold, average, and hot markets, showing how the role of concentration varies with the strength of the labor market.

TABLE XII: RECRUITING CONCENTRATION AND THE MINORITY EARNINGS-GAP RESPONSE TO TIGHTNESS

Outcome: W - M earnings gap	Estimate
Panel A. Linear specification	
Population: Beginning-of-quarter jobs	
Tightness	-0.21^{***}
	(0.05)
Tightness × concentration	-0.02
	(0.03)
Population: New stable hires	
Tightness	-0.28^{***}
	(0.09)
Tightness × concentration	-0.06
	(0.05)

Panel B. Earnings-gap responses by concentration and market temperature

	Overall (linear)	Cold U = 8.87%	Average U = 5.60%	Hot U = 2.93%
Population: Beginning-of-quarter jobs				
Tightness at p25 concentration	-0.19^{***}	-0.05	-0.27^{***}	-0.56^{***}
	(0.05)	(0.06)	(0.08)	(0.13)
Tightness at p75 concentration	-0.21^{***}	-0.15^{***}	-0.25^{***}	-0.38^{***}
	(0.05)	(0.06)	(0.08)	(0.12)
Significance of p25–p75 difference		†††		†
Population: New stable hires				
Tightness at p25 concentration	-0.22^{**}	-0.06	-0.41^{***}	-0.82^{***}
	(0.09)	(0.10)	(0.13)	(0.20)
Tightness at p75 concentration	-0.28^{***}	-0.19^{*}	-0.35^{***}	-0.56^{***}
	(0.09)	(0.10)	(0.13)	(0.19)
Significance of p25–p75 difference		†††		†

Notes: The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings. Responses to a 1-percentage-point unemployment decline, and differences between those responses, are expressed in percentage points of the earnings gap. All specifications include concentration, MSA-by-quarter-of-year and calendar-quarter fixed effects, and MSA-specific linear trends. Estimates are unweighted and use 20,176–20,326 observations in 368 MSAs over 2011Q3–2025Q2. Panel A reports the linear tightness response at median concentration and the percentage-point difference in that response associated with the p75–p25 concentration difference. Panel B evaluates responses at low and high concentration: the overall column uses the linear model, and the cold, average, and hot columns use a cubic in tightness interacted with concentration. Concentrations p25 and p75 use the occupational-reference sample and are approximately 942 and 2,680 on the zero-to-10,000 HHI scale. Appendix I.3 provides construction and reference values. MSA-clustered standard errors are in parentheses. Stars test individual responses: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Daggers test the p25–p75 difference: †, ††, and ††† denote the same respective thresholds.

Recruiting concentration limits the additional minority earnings-gap compression that tightening produces as labor markets heat up. The hot-market estimates in Panel B highlight this difference.

For beginning-of-quarter jobs, a 1-percentage-point unemployment decline narrows the minority earnings gap by an estimated 0.6 percentage points at low concentration, compared with 0.4 percentage points at high concentration. The same, but more pronounced, pattern holds for new stable hires.

As labor markets cool, the difference in the response to tightening between low- and high-concentration markets becomes smaller, and the ordering eventually reverses. For beginning-of-quarter jobs, a 1-percentage-point unemployment decline narrows the minority earnings gap by approximately 0.3 percentage points at average conditions at both concentration levels. In cold markets, the corresponding response is statistically indistinguishable from zero at low concentration and indicates a small reduction of approximately 0.1 percentage points at high concentration.

The offsetting concentration differences in cold and hot markets help explain why the linear estimates in Panel A show little overall difference in earnings compression. The estimated interactions are small and statistically indistinguishable from zero for both earnings populations. Together, these results suggest that employer competition shapes how strongly labor scarcity translates into relative earnings gains for minorities, particularly in already-tight markets.

10 Interpretation

We interpret the paper's findings through three major families of explanations linking labor-market conditions to earnings inequality between minorities and whites: differences in the kinds of jobs workers hold, differences in how employers evaluate workers and set the terms of a match, and differences in the set of jobs and employers workers can effectively reach.

The first, labor-market segmentation, locates the difference in the kinds of jobs workers hold. White and minority workers hold different jobs for reasons that can precede the contemporaneous hiring decision, including premarket skills and credentials (Neal and Johnson 1996), occupational and industry sorting, and the jobs available where workers live (Altonji and Blank 1999; Bayer and Charles 2018). Under this account, the gap can narrow as markets tighten if demand strengthens disproportionately in segments where minority workers are concentrated (Bound and Holzer 2000).

The second locates the difference between minorities and whites in how employers evaluate workers and set the terms of a match once a potential worker–employer match is under consideration. Employers may evaluate minority applicants differently because of preferences (Becker 1957) or

beliefs about characteristics they cannot observe (Lang and Lehmann 2012). Such discriminatory treatment can operate both in hiring decisions and in pay-setting. At the hiring stage, Kline et al. (2024) document lower callback rates for Black applicants than for otherwise comparable white applicants, with substantial differences across employers. The consequences of discriminatory screening can change with labor-market conditions.[26] As applicants become scarce, passing over suitable workers becomes more costly because employers have fewer alternatives for filling vacancies. For instance, the framework of Lang et al. (2005) illustrates how shortening applicant queues can reduce the earnings penalty from discriminatory hiring even when employers' racial preferences remain unchanged. Discriminatory treatment can also affect the pay workers receive. Using personnel records from a large service organization, Castilla (2008) finds evidence of discrimination in salary increases for minority employees with the same performance ratings as white employees, and Black (1995) shows how just a few prejudiced employers can lower minority workers' equilibrium wages even at unprejudiced firms. Tighter markets can make such pay disadvantages harder to sustain when they provide workers with better outside offers, increasing the risk that underpaid employees leave or reject unfavorable terms. The competing-offer mechanism in Cahuc et al. (2006) illustrates how this pressure can affect compensation: outside offers induce employers to renegotiate pay to recruit or retain workers.

The third locates the difference between minorities and whites in the set of jobs and employers workers can effectively reach. Information about vacancies and pay, referrals, networks, geography, and other search frictions shape that opportunity set (Bayer et al. 2008; Hellerstein et al. 2011; Montgomery 1991). If these channels differ across groups, workers can face different sets of feasible opportunities before any employer evaluates them. Tightening may therefore narrow differences between whites and minorities in effective opportunities if employers search more broadly or workers receive more offers; those additional opportunities can also strengthen the outside options of workers who remain in their current jobs.

All three explanations can coexist in equilibrium. For instance, networks can shape the opportunities available to workers, affecting their outside options and the discrimination they face at their current employers. Similarly, discrimination or unequal access earlier in workers' careers

[26] An example of how hiring standards vary with labor-market conditions is provided by Modestino et al. (2020), who show that employers raise education and experience requirements when workers are plentiful. Although their analysis does not examine race-specific differences, changes in these requirements can affect minority and white applicants differently because the two groups differ in the observed characteristics employers use to screen applicants.

can help shape the kinds of jobs they ultimately hold.[27] Our data do not allow us to separate these interconnected mechanisms.

Our results rule out particular channels through which labor-market segmentation could account for earnings compression, but do not allow us to distinguish discrimination from information and network explanations. Observed changes in workforce composition and higher labor demand in minority-concentrated occupations do not explain the effects of tightness on the earnings gap. Our results also provide limited evidence of systematically greater minority upgrading across broad job categories. These findings restrict explanations based on changes in who is employed or in demand and reallocation across broad job categories, while leaving open earnings improvements within those categories.

Discrimination and information or network barriers are harder to distinguish because both narrow the set of outside options available to minority workers, and tightening reduces the earnings consequences of either by widening those options. Employers may become more willing to hire minority applicants as vacancies become harder to fill. Tightening may also generate more offers outside personal networks, reducing the disadvantage of holding fewer or less valuable contacts for finding well-paying jobs. Neither employers' racial preferences nor workers' underlying information and network disadvantages need to change: more frequent offers alone can reduce their earnings consequences. Similar improvements in hiring and earnings can therefore arise under either explanation.

What our findings do indicate, however, is that competition among employers helps determine how strongly minority earnings gaps respond to tightening. The concentration results provide the clearest evidence: earnings compression strengthens as markets heat up, but the additional compression is smaller where recruiting is more concentrated. This pattern suggests that the availability of competing employers, not simply the overall strength of labor demand, shapes how strongly greater labor scarcity translates into pressure to improve pay.

The hiring and mobility analyses show that group differences in transition rates or destinations do not account for most of the observed minority earnings gap compression. This directs attention to improvements in earnings within the broad job categories we observe, which need not occur within the same employer. More competing offers can allow workers to move to better-paying em-

[27] For the purpose of distinguishing the three explanations, however, we treat observed differences in the kinds of jobs workers hold as given. This organizational distinction does not imply that those differences are independent of discrimination or unequal access.

ployers within those categories; alternatively, current employers may raise pay to retain workers who receive outside offers (Cahuc et al. 2006). The evidence therefore points to a role for employer competition in changing the earnings consequences of existing disadvantages, without identifying whether those disadvantages originated in job segmentation, discriminatory treatment, or unequal access to opportunities.

11 Conclusion

This paper shows that tighter labor markets narrow minority earnings gaps, with stronger compression as slack declines and particularly pronounced gains among recent hires. These findings show that the benefits of strong labor markets for minority workers extend beyond access to employment to relative earnings. The pattern is not readily explained by observed changes in workforce composition or job transitions. On the other hand, employer competition shapes the gains from tightening, particularly when labor is already scarce. Taken together, the results point to the importance of both strong labor demand and workers' access to competing employers for reducing the earnings disparities between minorities and whites.

An important open question is whether the relative earnings gains achieved in tight markets persist when demand weakens. Following workers beyond periods of strong demand would help distinguish temporary earnings compression from lasting improvements in career prospects. Another is which barriers tighter markets help workers overcome. Combining evidence on workers' access to job offers with direct evidence on employers' hiring and pay-setting decisions would help distinguish the roles of unequal access to opportunities and discriminatory treatment. Understanding both the mechanisms and the persistence of these gains would help clarify how strong labor demand and policies that broaden workers' opportunities can jointly contribute to lasting reductions in disparities between minority and white workers.

References

Aaronson, D., Hu, L., & Rajan, A. (2020). How much did the minimum wage drive real wage growth during the late 2010s? *Chicago Fed Letter*, (435).

Aaronson, S. R., Daly, M. C., Wascher, W. L., & Wilcox, D. W. (2019). Okun revisited: Who benefits most from a strong economy? *Brookings Papers on Economic Activity*, *50*(1), 333–404.

Abdelwahed, L. (2025). *The rise and fall of DEI in corporate America*. Revelio Labs. Retrieved 2026, from https://www.reveliolabs.com/news/social/the-rise-and-fall-of-dei-in-corporate-america

Akerlof, G. A., Rose, A. K., & Yellen, J. L. (1988). Job switching and job satisfaction in the U.S. labor market. *Brookings Papers on Economic Activity*, *19*(2), 495–594.

Allegretto, S. A., Dube, A., & Reich, M. (2011). Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data. *Industrial Relations*, *50*(2), 205–240. https://doi.org/10.1111/j.1468-232X.2011.00634.x

Altonji, J. G., & Blank, R. M. (1999). Race and gender in the labor market. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (pp. 3143–3259, Vol. 3). Elsevier. https://doi.org/10.1016/S1573-4463(99)30039-0

Autor, D., Dube, A., & McGrew, A. (2024). *The unexpected compression: Competition at work in the low wage labor market* (Working Paper No. 31010) (Revised version). National Bureau of Economic Research.

Ayas, R., Tilly, P., & Rawlings, D. (2023). *Cutting costs at the expense of diversity*. Revelio Labs. Retrieved 2026, from https://www.reveliolabs.com/news/social/cutting-costs-at-the-expense-of-diversity

Azar, J., Marinescu, I., Steinbaum, M., & Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, *66*, Article 101886. https://doi.org/10.1016/j.labeco.2020.101886

Başkaya, Y. S., & Rubinstein, Y. (2012). *Using federal minimum wages to identify the impact of minimum wages on employment and earnings across the U.S. states* [Working paper, December 2012 version].

Bayer, P., Ross, S. L., & Topa, G. (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy*, *116*(6), 1150–1196. https://doi.org/10.1086/595975

Bayer, P., & Charles, K. K. (2018). Divergent paths: A new perspective on earnings differences between Black and white men since 1940. *The Quarterly Journal of Economics*, *133*(3), 1459–1501. https://doi.org/10.1093/qje/qjy003

Becker, G. S. (1957). *The economics of discrimination*. University of Chicago Press.

Benigno, P., & Eggertsson, G. B. (2023). *It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve* (NBER Working Paper No. 31197). National Bureau of Economic Research.

Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of Labor Economics*, *13*(2), 309–334. https://doi.org/10.1086/298376

Bound, J., & Holzer, H. J. (2000). Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, *18*(1), 20–54. https://doi.org/10.1086/209949

Burya, A., Mano, R., Timmer, Y., & Weber, A. (2023). The wage Phillips curve under labor market power. *AEA Papers and Proceedings*, *113*, 110–113. https://doi.org/10.1257/pandp.20231007

Cahuc, P., Postel-Vinay, F., & Robin, J.-M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, *74*(2), 323–364. https://doi.org/10.1111/j.1468-0262.2006.00665.x

Cajner, T., Radler, T., Ratner, D., & Vidangos, I. (2017). *Racial gaps in labor market outcomes in the last four decades and over the business cycle* (Finance and Economics Discussion Series No. 2017-071). Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2017.071

Castilla, E. J. (2008). Gender, race, and meritocracy in organizational careers. *American Journal of Sociology*, *113*(6), 1479–1526. https://doi.org/10.1086/588738

Cavounidis, C., Lang, K., & Weinstein, R. (2024). The boss is watching: How monitoring decisions hurt black workers. *The Economic Journal*, *134*(658), 485–514. https://doi.org/10.1093/ej/uead079

Choi, H., & Marinescu, I. (2024). Data for labor market concentration using Lightcast (formerly Burning Glass Technologies). *Data in Brief*, *55*, Article 110647. https://doi.org/10.1016/j.dib.2024.110647

Couch, K. A., & Fairlie, R. (2010). Last hired, first fired? Black–white unemployment and the business cycle. *Demography*, *47*(1), 227–247. https://doi.org/10.1353/dem.0.0086

Daly, M. C., Hobijn, B., & Pedtke, J. H. (2017). *Disappointing facts about the Black-white wage gap* (FRBSF *Economic Letter No. 2017-26*). Federal Reserve Bank of San Francisco.

Daly, M. C., Hobijn, B., & Pedtke, J. H. (2020). Labor market dynamics and Black–white earnings gaps. *Economics Letters*, *186*, Article 108807. https://doi.org/10.1016/j.econlet.2019.108807

Diamond, R. (2016). The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *American Economic Review*, *106*(3), 479–524. https://doi.org/10.1257/aer.20131706

Domash, A., & Summers, L. H. (2022). *How tight are U.S. labor markets?* (NBER Working Paper No. 29739). National Bureau of Economic Research.

Dube, A., Lester, T. W., & Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *The Review of Economics and Statistics*, *92*(4), 945–964. https://doi.org/10.1162/REST_a_00039

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Backman, D., Breton, E., Cooper, G., Rivera Drew, J. A., Richards, S., & Van Riper, D. (2025). *Integrated public use microdata series, current population survey* (Version 13.0). Minneapolis, MN, IPUMS. https://doi.org/10.18128/D030.V13.0

Ganong, P., & Shoag, D. (2017). Why has regional income convergence in the U.S. declined? *Journal of Urban Economics*, *102*, 76–90. https://doi.org/10.1016/j.jue.2017.07.002

Gerard, F., Lagos, L., Severnini, E., & Card, D. (2021). Assortative matching or exclusionary hiring? The impact of employment and pay policies on racial wage differences in Brazil. *American Economic Review*, *111*(10), 3418–3457. https://doi.org/10.1257/aer.20181596

Gertler, M., Huckfeldt, C., & Trigari, A. (2020). Unemployment fluctuations, match quality, and the wage cyclicality of new hires. *The Review of Economic Studies*, *87*(4), 1876–1914. https://doi.org/10.1093/restud/rdaa004

Haltiwanger, J. C., Hyatt, H. R., Kahn, L. B., & McEntarfer, E. (2018). Cyclical job ladders by firm size and firm wage. *American Economic Journal: Macroeconomics*, *10*(2), 52–85. https://doi.org/10.1257/mac.20150245

Hellerstein, J. K., McInerney, M., & Neumark, D. (2011). Neighbors and coworkers: The importance of residential labor market networks. *Journal of Labor Economics*, *29*(4), 659–695. https://doi.org/10.1086/660776

Hines, J. R., Jr., Hoynes, H. W., & Krueger, A. B. (2001). *Another look at whether a rising tide lifts all boats* (Working Paper No. 8412). National Bureau of Economic Research. https://doi.org/10.3386/w8412

Hoynes, H. W. (2000). The employment, earnings, and income of less skilled workers over the business cycle. In D. Card & R. M. Blank (Eds.), *Finding jobs: Work and welfare reform* (pp. 23–71). Russell Sage Foundation.

Indeed & Glassdoor. (2022). *Hiring and workplace trends report 2023*. Indeed and Glassdoor.

Johnson, M. S., Lavetti, K., & Lipsitz, M. (2025). The labor market effects of legal restrictions on worker mobility. *Journal of Political Economy*, *133*(9), 2735–2793. https://doi.org/10.1086/736217

Jones-Ruiz, K. (2020). *What began as a small startup is now a successful partnership between states and US Census Bureau*. US Census Bureau. Retrieved 2026, from https://www.census.gov/library/stories/2020/01/census-bureau-innovative-lehd-program-local-workforce-dynamics-turns-20.html

Kline, P., Rose, E. K., & Walters, C. R. (2024). A discrimination report card. *American Economic Review*, *114*(8), 2472–2525. https://doi.org/10.1257/aer.20230700

Lang, K., Manove, M., & Dickens, W. T. (2005). Racial discrimination in labor markets with posted wage offers. *American Economic Review*, *95*(4), 1327–1340. https://doi.org/10.1257/0002828054825547

Lang, K., & Lehmann, J.-Y. K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, *50*(4), 959–1006. https://doi.org/10.1257/jel.50.4.959

Langella, M., & Manning, A. (2021). Marshall Lecture 2020: The measure of monopsony. *Journal of the European Economic Association*, *19*(6), 2929–2957. https://doi.org/10.1093/jeea/jvab039

Lightcast. (2026). *Data collection and processing methodology*. Retrieved 2026, from https://kb.lightcast.io/en/articles/7124763-data-collection-and-processing-methodology

Management Leadership for Tomorrow. (2021). *MLT announces newest cohort of industry leading employers that will pursue pathbreaking MLT Black Equity at Work certification*. Retrieved 2026, from https://mlt.org/blog/news-story-old/black-equity-at-work-certification-newest-cohort/

Martins, P. S., Solon, G., & Thomas, J. P. (2012). Measuring what employers do about entry wages over the business cycle: A new approach. *American Economic Journal: Macroeconomics*, *4*(4), 36–55. https://doi.org/10.1257/mac.4.4.36

Meer, J., & West, J. (2016). Effects of the minimum wage on employment dynamics. *Journal of Human Resources*, *51*(2), 500–522. https://doi.org/10.3368/jhr.51.2.0414-6298R1

Modestino, A. S., Shoag, D., & Ballance, J. (2020). Upskilling: Do employers demand greater skill when workers are plentiful? *The Review of Economics and Statistics*, *102*(4), 793–805.

Montgomery, J. D. (1991). Social networks and labor-market outcomes: Toward an economic analysis. *American Economic Review*, *81*(5), 1408–1418.

Moscarini, G., & Postel-Vinay, F. (2018). The cyclical job ladder. *Annual Review of Economics*, *10*, 165–188. https://doi.org/10.1146/annurev-economics-080217-053425

Neal, D. A., & Johnson, W. R. (1996). The role of premarket factors in Black-white wage differences. *Journal of Political Economy*, *104*(5), 869–895. https://doi.org/10.1086/262045

Neumark, D., Salas, J. M. I., & Wascher, W. (2014). Revisiting the minimum wage–employment debate: Throwing out the baby with the bathwater? *ILR Review*, *67*(Supplement), 608–648. https://doi.org/10.1177/00197939140670S307

North Carolina Agricultural and Technical State University. (2024). *Office of career services: December 2024 newsletter*. Retrieved 2026, from https://www.ncat.edu/provost/career-services/newsletter/december-2024/index.php

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica*, *25*(100), 283–299.

Pissarides, C. A. (2000). *Equilibrium unemployment theory* (2nd ed.). MIT Press.

Ross, M., & Bateman, N. (2019). *Meet the low-wage workforce*. Brookings Institution.

Solon, G., Barsky, R., & Parker, J. A. (1994). Measuring the cyclicality of real wages: How important is composition bias? *The Quarterly Journal of Economics*, *109*(1), 1–25. https://doi.org/10.2307/2118426

Sorkin, I. (2018). Ranking firms using revealed preference. *The Quarterly Journal of Economics*, *133*(3), 1331–1393. https://doi.org/10.1093/qje/qjy001

Sorkin, I. (2025). *Quantifying racial disparities using consecutive employment spells*. Federal Reserve Bank of Minneapolis Institute, Working Paper 120. https://doi.org/10.21034/iwp.120

US Bureau of Labor Statistics. (2018). *Local area unemployment statistics: Concepts*. Handbook of Methods. Retrieved 2026, from https://www.bls.gov/opub/hom/lau/concepts.htm

US Bureau of Labor Statistics. (2026a). *Consumer price index* [CPI-U, US city average, all items, not seasonally adjusted, series CUUR0000SA0, retrieved through the BLS public data API]. Retrieved 2026, from https://www.bls.gov/cpi/

US Bureau of Labor Statistics. (2026b). *Job openings and labor turnover survey: Data definitions*. Retrieved 2026, from https://www.bls.gov/jlt/jltdef.htm

US Census Bureau. (2024). *Place of work for workers 16 years and over—metropolitan statistical area level* [Table B08016, United States. National estimates retrieved through the Census Reporter data API]. American Community Survey, 1-Year Estimates Detailed Tables. Retrieved 2026, from https://data.census.gov/table/ACSDT1Y2024.B08016

US Census Bureau. (2026a). *QWI measure descriptions*. Longitudinal Employer–Household Dynamics. Retrieved 2026, from https://lehd.ces.census.gov/doc/QWI_101.pdf

US Census Bureau. (2026b). *Job-to-job flows: New statistics on worker flows across jobs*. Longitudinal Employer–Household Dynamics. Retrieved 2026, from https://lehd.ces.census.gov/doc/J2J_quickstart_guide.pdf

US Census Bureau. (2026c). *American Community Survey: Public use microdata sample*. Retrieved 2026, from https://www.census.gov/programs-surveys/acs/microdata.html

Vaghul, K., & Zipperer, B. (2022). *Historical state and sub-state minimum wages* (1.4.0). Retrieved 2026, from https://github.com/benzipperer/historicalminwage/releases/tag/v1.4.0

Wursten, J., & Reich, M. (2023). Racial inequality in frictional labor markets: Evidence from minimum wages. *Labour Economics*, *82*, Article 102344. https://doi.org/10.1016/j.labeco.2023.102344

Online Appendix (List of Contents)

C Data **51**

C.1 QWI Coverage and Earnings Construction . . . 51

C.2 Household Surveys and Earnings Comparisons . . . 52

C.3 Unemployment, Postings, and Job Openings . . . 54

D Empirical Framework **56**

D.1 Composition Controls and Local Linear Trends . . . 56

E Core Results **59**

E.1 Employment-Weighted Estimates . . . 59

E.2 State Dependence Excluding 2020 . . . 61

E.3 Unemployment and Recruiting Intensity Entered Together . . . 63

E.4 Online Postings against the Official Vacancy Count . . . 64

F Assessing the Role of Composition and Policy **66**

F.1 Worker-side Composition . . . 66

F.1.1 Workforce and Population Controls . . . 66

F.1.2 Cross-MSA Flows and Earnings Accounting . . . 68

F.2 Vacancy-side Composition: Occupational Tightness Construction . . . 70

F.3 Policy Levers . . . 71

F.3.1 Minimum-Wage Construction and Comparison . . . 71

F.3.2 Diversity, Equity, and Inclusion . . . 72

G Transition Frequency and Earnings-Gap Compression **80**

G.1 Transition Rates from Job-to-Job Flows . . . 80

G.2 Initial and Stable Hiring . . . 81

G.3 Flow Controls in the Earnings Regression . . . 81

H Transition Quality and Earnings-Gap Compression **82**

H.1 Destination Pay Ranks . . . 82

H.2 Separation between Whites and Minorities in Hiring . . . 83

H.3 Earnings Gains across Transitions . . . 83

I Monopsony Power **85**

I.1 Measuring Recruiting Concentration . . . 85

I.2 Benchmarking the Measure . . . 86

I.3 Specifications and Reported Responses . . . 87

C Data

C.1 QWI Coverage and Earnings Construction

We use release R2026Q2 of the QWI, retaining the race-by-ethnicity tabulations for private-sector and state and local government employment. The main analysis uses all-industry totals. For MSAs that cross state boundaries, the QWI report each state's portion separately. We combine these portions by weighting their mean earnings by the corresponding job counts: beginning-of-quarter jobs for beginning-of-quarter earnings, and new stable hires for new-stable-hire earnings. An MSA mean is available only when every reported state portion has usable earnings and a matching job count.

The beginning-of-quarter outcome uses `EarnBeg` and its corresponding count, `Emp`. A job is counted in quarter t when earnings from the same employer are positive in $t-1$ and t. The new-stable-hire outcome uses `EarnHirNS` and `HirNS`: earnings from the employer must be positive in $t-1$, t, and $t+1$, and absent in $t-2$ through $t-5$. Entry thus occurs in $t-1$, while earnings are measured in t. These conditions identify a new employment relationship that reaches full-quarter duration and exclude recalls within the preceding year. They do not establish continuous hours worked or full-time employment within the quarter. Both earnings variables equal the corresponding group's quarterly earnings per job divided by three.

All MSA-level analyses exclude the same 13 metropolitan areas with fewer than 100 white or pooled Black-and-Hispanic new stable hires in at least one quarter of 2010Q1–2025Q2. The exclusion applies to both earnings populations in every period, including the pre-2010 historical-period regressions. It leaves 374 of the original 387 MSAs in the main panel; individual exercises impose additional data requirements. National and state-level descriptive comparisons retain their original geographic coverage.

Across the 374 retained MSAs and 60 quarters of 2010Q3–2025Q2, earnings gaps are available for 98.7 percent of possible MSA–quarters for beginning-of-quarter jobs and 98.0 percent for new stable hires. Most missing MSA–quarters occur toward the end of the sample, as the Alaska and Michigan series in our panel end in 2016Q2 and 2021Q4, respectively.

White and Black workers are non-Hispanic; Hispanic workers may be of any race. Minority mean earnings are the weighted average of Black and Hispanic mean earnings, using each group's share of minority jobs as its weight. The white-minus-minority earnings gap is 100 times the difference

between log white mean earnings and log pooled minority mean earnings, and it is constructed only when all three group means and matching job counts are available.

C.2 Household Surveys and Earnings Comparisons

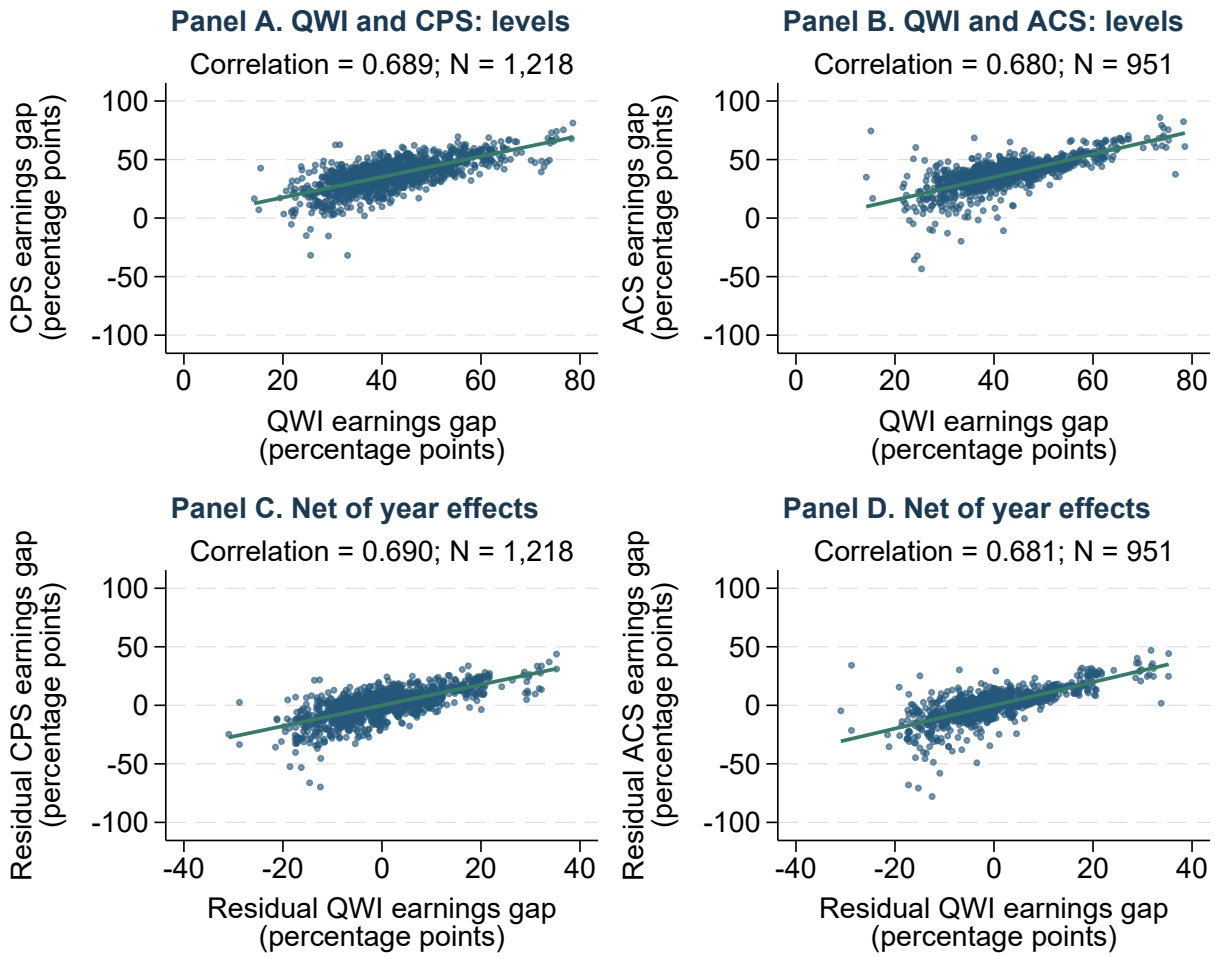
The CPS earnings gap comparison discussed in Section 3 and shown in Figure C.1 of this appendix uses outgoing-rotation-group respondents in the basic monthly CPS, provided through IPUMS (Flood et al. 2025). We retain employed civilian nonstudents ages 16–65 in private-sector or state and local government wage-and-salary jobs, with positive usual weekly earnings and earnings weights. We exclude allocated earnings, federal employees, and the self-employed. We multiply observations at the old fixed earnings top code by 1.5, without applying this adjustment to the newer dynamic top-coded values. The weekly-earnings sample does not require a valid hourly wage. Annual state means pool eligible observations across all twelve months using earnings weights.

The equivalent ACS comparison uses the Census Bureau's one-year public-use microdata (US Census Bureau 2026c). We retain civilian household residents ages 16–64 who are not students, report private-sector or state and local government wage-and-salary work, and have positive wage income, weeks worked, and usual weekly hours, with no positive self-employment income. Respondents need not be employed at the interview because earnings refer to the preceding twelve months. We construct a weekly earnings measure for the ACS comparison of earnings gaps with the QWI data. Weekly earnings equal annual wage and salary income, adjusted by the Census income-adjustment factor, divided by weeks worked. We use reported weeks in 2005–2007 and from 2019 onward, and the midpoints of the reported weeks categories in 2008–2018. Annual group means use person weights. Standard one-year ACS data are unavailable for 2020.

Figure C.1 compares state-level minority earnings gaps in the QWI with those in the CPS and ACS to assess whether the sources capture similar geographic patterns. The figure shows unadjusted gaps and gaps net of year fixed effects. Each observation covers a full state, including its nonmetropolitan areas, rather than averaging MSA gaps within the state. Each QWI annual group mean weights quarterly `EarnBeg` by the matching `Emp` counts, requiring usable observations in all four quarters for whites, Blacks, and Hispanics. In every source, Black and Hispanic earnings totals and weights are pooled before calculating the minority mean.

The QWI and household surveys show similar geographic patterns in minority earnings gaps,

FIGURE C.1: MINORITY EARNINGS GAPS IN ADMINISTRATIVE AND HOUSEHOLD-SURVEY DATA



Notes: Each observation is a state–year, including metropolitan and nonmetropolitan areas; the District of Columbia is treated as a state. The white-minus-minority earnings gap is 100 times the difference in log mean earnings. Both axes are measured in percentage points. Minorities pool non-Hispanic Black workers and Hispanic workers of any race; white workers are non-Hispanic. The QWI measure beginning-of-quarter monthly earnings in private and state and local government jobs at all ages. Its annual means weight quarterly earnings by matching job counts. The CPS uses earnings-weighted usual weekly earnings for employed civilian nonstudents ages 16–65 in the same ownership categories, excluding allocated earnings and applying the fixed-top-code adjustment. The ACS uses person-weighted annual wage income divided by weeks worked for civilian household nonstudents ages 16–64 with positive wage income, weeks, and hours and no positive self-employment income. Appendix C.2 gives the detailed eligibility rules. Minority earnings totals and weights are pooled before taking logs in each source. The QWI–CPS comparison covers 2000–2024; QWI–ACS covers 2005–2024 except 2020, when standard one-year ACS data are unavailable. Each pair requires complete years in both sources; its levels and residual panels use identical observations. Residuals remove year fixed effects separately in each source, with no state fixed effects, trends, or other controls. Each state–year receives equal weight in the Pearson correlation and linear fit. No observations are trimmed and no additional minimum survey-record count is imposed. Vertical axes display -100 to 100 percentage points of the earnings gap; any observations outside this range remain in the fitted lines, correlations, and sample counts.

both in levels and after removing common annual movements. The correlation between QWI and CPS gaps is approximately 0.69 before and after removing year fixed effects; the corresponding correlations with the ACS are approximately 0.68.

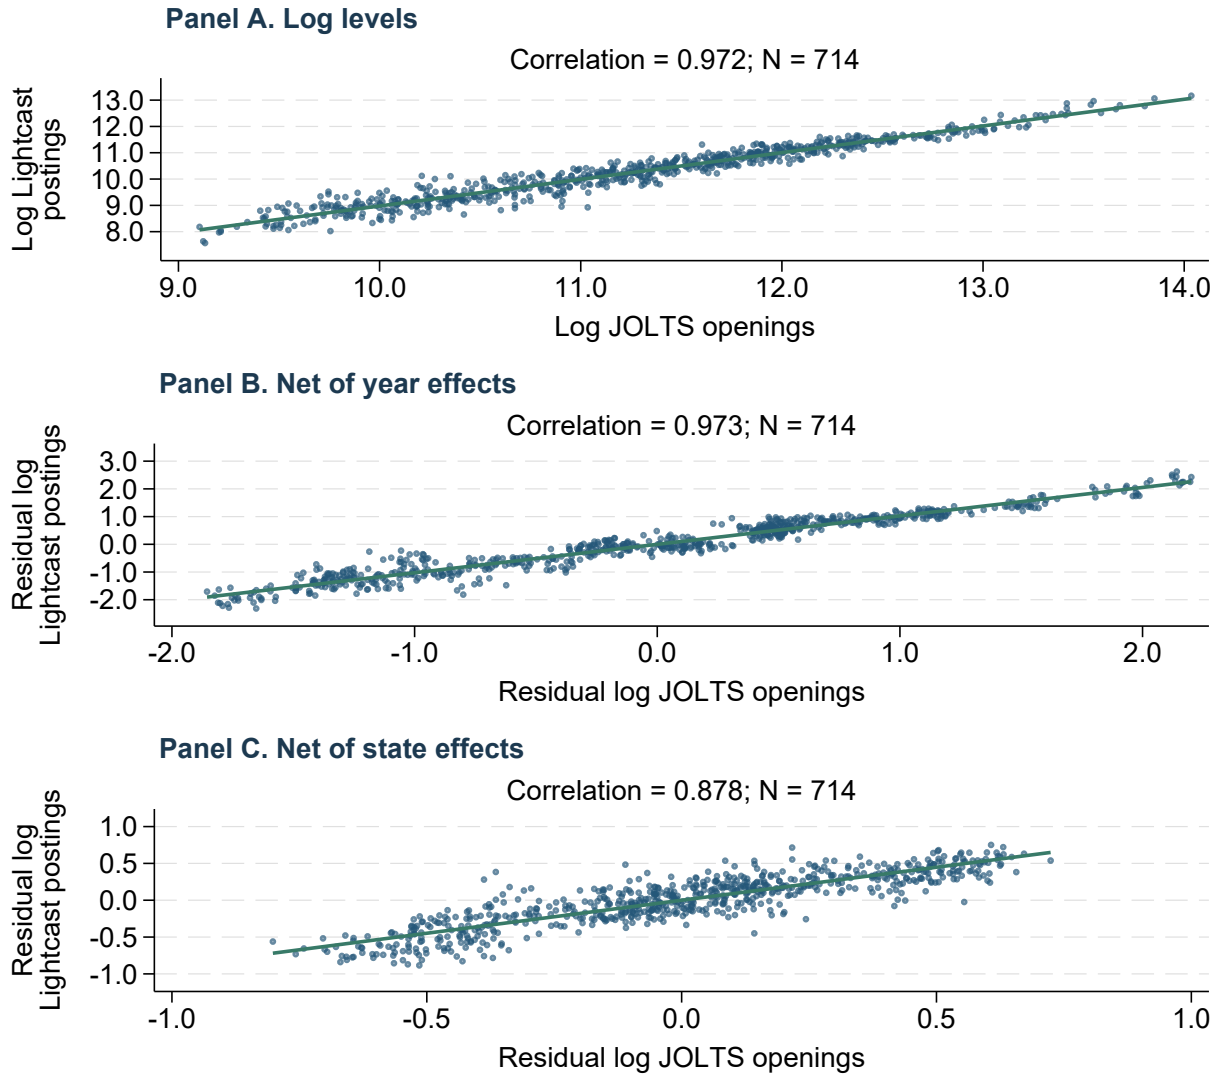
C.3 Unemployment, Postings, and Job Openings

Quarterly MSA unemployment rates are averages of the three monthly, not-seasonally-adjusted LAUS rates. The postings-to-unemployment measure divides the mean monthly number of newly dated Lightcast advertisements by the mean monthly number of unemployed people in the same quarter, then takes the natural log. Thus, it is the log of a ratio of quarterly averages, not an average of monthly ratios or logs. We take posting locations from Lightcast's county variable and map counties to MSAs using the 2023 metropolitan delineations. The local recruiting measure requires usable county and posting-date information, available for 99.1 percent of advertisements over 2010Q3–2025Q2. The share of postings with usable county and date information was lowest in 2023, at 95 percent, and recovered thereafter.

For the comparison with JOLTS, we aggregate Lightcast advertisements with identified counties to states and use the BLS's not-seasonally-adjusted total-nonfarm state job-opening counts. JOLTS openings are positions available on the last business day of the month for which employers are actively recruiting externally (US Bureau of Labor Statistics 2026b). We average all twelve monthly counts in each source and then take logs. Figure C.2 compares the paired state–year observations in log levels (Panel A), after removing year fixed effects (Panel B), and after removing state fixed effects (Panel C).

Lightcast postings and JOLTS openings track each other closely across states and over time. Their correlation is 0.972 in log levels and 0.973 after removing year fixed effects. After removing persistent differences across states, the correlation remains high, at 0.878, indicating that the two sources also capture similar year-to-year changes within states.

FIGURE C.2: STATE RECRUITING ACTIVITY IN JOLTS AND LIGHTCAST



Notes: All three panels use the same 714 state–years: the 50 states and the District of Columbia in 2011–2024, including 2020. JOLTS measures not-seasonally-adjusted total-nonfarm positions open on the last business day of the month for which employers are actively recruiting externally. Lightcast measures newly dated online job advertisements with an identified US county, summed across industries, including unclassified advertisements. Each annual count is the average of all twelve monthly counts; we take its natural log after averaging. Counts are measured in individual openings or advertisements, not thousands. Requiring complete years within the main 2010Q3–2025Q2 window yields the stated annual coverage. Panel A places log JOLTS openings on the horizontal axis and log Lightcast postings on the vertical axis. Panel B removes year fixed effects separately from each logged series. Panel C instead removes state fixed effects, retaining year-to-year variation. No state-specific trends or other controls are included. Green lines show linear fits; the reported Pearson correlations and fits give equal weight to each state–year and retain all observations.

D Empirical Framework

D.1 Composition Controls and Local Linear Trends

Table D.1 compares how direct composition controls and two alternative adjustments for local change affect the estimated earnings-gap response to tightness. Panel A considers MSA-specific linear trends, while Panel B considers state-by-year fixed effects. Both panels start from the same specification with MSA-by-quarter-of-year and calendar-quarter fixed effects. The former allow each MSA its own recurring seasonal pattern, with quarters of the year numbered one through four; the latter absorb shocks common to all MSAs in each calendar quarter. Each panel reports the share of this starting estimate absorbed by composition controls, the alternative adjustment, both together, and the increment from adding composition controls to that adjustment. Comparing the fourth column across panels shows whether composition controls have less influence on the tightness estimate after local trends or after state-level annual changes are absorbed. Every specification uses the same observations within each job population, so differences reflect the adjustment and not the sample.

The composition measures are mean age, the share holding a bachelor's degree or more, the foreign-born share, and the pooled Black and Hispanic share. We use MSA-wide measures rather than differences between minorities and whites because the exercise assesses whether local trends account for demographic changes relevant to the estimated earnings-gap response to tightness. Such changes may affect the two groups differently even when measured for the local population as a whole. Each is a person-weighted mean over civilian household residents aged 16–64 in the ACS one-year person files, excluding residents of group quarters and members of the Armed Forces. Residential PUMAs are assigned to their majority MSA under the 2023 metropolitan delineations, using the crosswalk appropriate to each PUMA vintage, assigning only those that lie predominantly within a single MSA. The measures describe the local working-age population as a whole, in work or out of it, and cover residents beyond the QWI-covered jobs that form the dependent variable.

Each annual value enters lagged one year, so that the composition attached to a quarter precedes the earnings and tightness measured in it. The Census Bureau released no standard one-year file for 2020, and the four quarters of 2021 therefore carry the 2019 values; no 2020 composition observation is imputed. The majority-PUMA assignment, which retains only MSAs for which the

combined omission and commission error is below 15 percent in both PUMA vintages, reaches 237 of the 374 metropolitan areas retained in the main panel. The first usable control year is 2012, so the estimation window runs 2013Q1–2025Q2 and retains 2020. Both panels contain 11,674 MSA–quarters for beginning-of-quarter jobs and 11,652 for new stable hires, in 237 MSAs, of which 944 and 943 observations fall in 2020. One new-stable-hire observation that is a singleton under state-by-year fixed effects is excluded from every specification to maintain a common sample across the two panels.

In Panel A, composition controls absorb 17.6 percent of the starting fixed-effects-only estimate for beginning-of-quarter jobs and 14.1 percent for new stable hires. Adding local trends instead absorbs 19.4 and 16.9 percent, respectively. Once trends are included, adding composition controls changes the estimates by only -1.6 and 1.2 percent of the same starting estimates, respectively. The exercise therefore indicates that the trends absorb the drift these four measures record, and that direct controls have little left to remove once they are imposed. In Panel B, adding composition controls to state-by-year fixed effects changes the estimates by amounts equivalent to 11.3 percent of the starting fixed-effects-only estimate for beginning-of-quarter jobs and 10.9 percent for new stable hires. Their larger additional influence under state-by-year fixed effects supports using MSA-specific trends to account for the measured demographic changes relevant to the tightness estimate. It does not establish that the trends absorb every source of slow-moving confounding, and the four measures are not an exhaustive description of local compositional change.

TABLE D.1: LOCAL LINEAR TRENDS AND MEASURED COMPOSITIONAL DRIFT

Panel A. MSA-specific linear trends

Job population	Composition controls	MSA trends	Both	Composition added to trends
Beginning-of-quarter jobs	17.6	19.4	17.7	-1.6
New stable hires	14.1	16.9	18.1	1.2

Panel B. State-by-year fixed effects

Job population	Composition controls	State-by-year fixed effects	Both	Composition added to state-by-year FE
Beginning-of-quarter jobs	17.6	-58.4	-47.1	11.3
New stable hires	14.1	-1.5	9.4	10.9

Notes: Entries report the share of the same starting tightness estimate absorbed by each adjustment, in percent. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean earnings, and tightness is the negative LAUS unemployment rate. Both panels start from MSA-by-quarter-of-year and calendar-quarter fixed effects alone, with estimates of -0.43 (standard error 0.11) for beginning-of-quarter jobs and -0.45 (0.10) for new stable hires, in percentage points per 1-percentage-point unemployment decline. Panel A adds composition controls, MSA-specific linear trends, and both; its fourth column reports the additional share absorbed when composition controls are added to trends. Panel B starts from the same baseline and adds composition controls, state-by-year fixed effects, and both; its fourth column reports the additional share absorbed when composition controls are added to state-by-year fixed effects. No specification in Panel B includes MSA trends. All entries, including the fourth column, use the common starting estimate as the denominator; a negative entry means the estimate grows in magnitude. The controls are mean age, the share with a bachelor's degree or more, the foreign-born share, and the minority share. They use person-weighted ACS one-year data for civilian household residents ages 16–64, assigned to MSAs by majority residential PUMA under the 2023 delineations. Controls enter lagged one year; 2019 supplies the controls for 2021 because no standard one-year ACS file is available for 2020. The sample covers 2013Q1–2025Q2, including 2020, with 11,674 beginning-of-quarter and 11,652 new-stable-hire observations in 237 MSAs. All specifications in both panels use identical observations within each population, excluding one new-stable-hire singleton throughout to maintain this common sample. MSA–quarters are equally weighted, and standard errors are clustered by MSA. All MSA regressions exclude, in every period, the 13 metropolitan areas with fewer than 100 white or minority new stable hires in at least one quarter of 2010Q1–2025Q2.

E Core Results

E.1 Employment-Weighted Estimates

Table E.1 repeats the specifications of Table I under the alternative weighting. Each MSA–quarter is weighted by its minority job count in the relevant job population, and the five largest metropolitan areas are excluded. Everything else follows the main table.

The two weightings answer different questions. Equal weights describe the response of the average metropolitan labor market, while minority job counts give greater influence to markets with more minority employment. Excluding New York, Los Angeles, Chicago, Miami, and Houston limits the influence of the very largest markets. Even after these exclusions, the 20 largest remaining markets receive about 45 percent of the employment weight, compared with about 5 percent if each retained MSA–quarter is weighted equally. The 200 smallest receive only about 8–9 percent, rather than about 54 percent. Therefore, we report the equal-weighted estimates in the main text because the paper's question concerns how minority earnings gaps respond to tightening across local labor markets.

Under this alternative weighting and sample, the point estimates in our preferred specification remain close to those in the main table.

TABLE E.1: TIGHTNESS AND THE MINORITY EARNINGS GAP, EMPLOYMENT-WEIGHTED

	$-U$			$\ln(V/U)$		
Outcome: W - M earnings gap	MSA seasonal FE only	Add state-quarter FE	Add date FE and trends	MSA seasonal FE only	Add state-quarter FE	Add date FE and trends
Population: Beginning-of-quarter jobs						
Tightness	-0.62^{***}	-1.20^{***}	-0.21^{***}	-1.91^{***}	-1.54^{***}	-0.25
	(0.07)	(0.35)	(0.08)	(0.20)	(0.56)	(0.17)
Population: New stable hires						
Tightness	-1.28^{***}	-0.78^{***}	-0.30^{***}	-3.69^{***}	-0.85^{**}	-0.44^{**}
	(0.09)	(0.25)	(0.11)	(0.20)	(0.40)	(0.22)

Notes: This table mirrors Table I, weighting each MSA–quarter by minority job counts and excluding the five largest metropolitan areas. The outcome is the white-minus-minority earnings gap in the indicated job population, constructed as 100 times the difference in log mean QWI monthly earnings. Each cell reports the estimated response of the gap in percentage points from equation (1), scaled to the changes in tightness described below. Tightness is the LAUS unemployment rate, which carries a negative sign so that the measure rises as a market tightens, or the log ratio of Lightcast job postings to LAUS unemployment. Entries for the first are read for a 1-percentage-point decline in unemployment and for the second for a doubling of the ratio. Standard errors are clustered by MSA. The estimation sample covers 2010Q3–2025Q2, including 2020, with the all-period MSA exclusions described in Table D.1. The weighted samples contain 369 MSAs. The state-by-calendar-quarter specifications drop 360 singleton observations for beginning-of-quarter jobs and 377 for new stable hires, leaving 363 MSAs in each case. Every specification includes MSA-by-quarter-of-year fixed effects. The state-by-calendar-quarter columns add those effects; the trend columns instead add calendar-quarter effects and MSA-specific linear trends. Job populations are defined in Section 3. Parentheses report standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E.2 State Dependence Excluding 2020

Table E.2 repeats the state-dependence exercise of Table II on the sample that excludes the four quarters of 2020. We estimate the same cubic specification in negative unemployment. As in the main table, cold and hot conditions use the 10^{th} and 90^{th} percentiles of tightness outside 2020, while average conditions use the mean over the full sample. Panel A reports the fitted slope at unemployment rates in cold, average, and hot markets. Panel B reports the differences between those slopes. Because the evaluation points are held at the values used in the main table, the comparison isolates the change in estimation sample rather than combining it with a change in where the polynomial is read.

TABLE E.2: THE EFFECT OF TIGHTNESS AT COLD, AVERAGE AND HOT MARKET CONDITIONS, EXCLUDING 2020

Outcome: W - M earnings gap	Cold U = 8.87%	Difference: avg. – cold	Average U = 5.60%	Difference: hot – avg.	Hot U = 2.93%	Difference: hot – cold
Panel A. Tightness slopes at each market condition						
Population: Beginning-of-quarter jobs						
Tightness	−0.11	—	−0.25***	—	−0.39***	—
	(0.08)		(0.09)		(0.14)	
Population: New stable hires						
Tightness	−0.20*	—	−0.40***	—	−0.67***	—
	(0.11)		(0.13)		(0.21)	
Panel B. Differences between the tightness slopes						
Population: Beginning-of-quarter jobs						
Tightness	—	−0.14**	—	−0.13	—	−0.28*
		(0.07)		(0.08)		(0.15)
Population: New stable hires						
Tightness	—	−0.20**	—	−0.28**	—	−0.47**
		(0.09)		(0.11)		(0.20)

Notes: This table repeats Table II excluding 2020. The outcome is the white-minus-minority earnings gap, constructed as 100 times the difference in log mean earnings, and tightness is the negative LAUS unemployment rate. Entries are in percentage points. Panel A reports local responses to a 1-percentage-point unemployment decline; Panel B reports differences between these responses, with standard errors using their covariance within each cubic specification. Evaluation points remain fixed at unemployment rates of 8.87, 5.60, and 2.93 percent to isolate the change in sample. All regressions include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends, weight MSA–quarters equally, and cluster standard errors by MSA. The sample covers 2010Q3–2025Q2, excluding the four quarters of 2020, with 20,670 beginning-of-quarter and 20,507 new-stable-hire observations in 374 MSAs. All-period MSA exclusions follow Table D.1. Parentheses report standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Excluding 2020 increases the difference between cold- and hot-market responses for both job populations. The cold-to-hot increase in earnings compression rises from 0.16 to 0.28 percentage points for beginning-of-quarter jobs and from 0.41 to 0.47 percentage points for new stable hires,

per 1-percentage-point unemployment decline. These contrasts are significant at the 10 and 5 percent levels, respectively, when 2020 is excluded. The pandemic quarters therefore do not generate the estimated increase in compression as markets heat up.

E.3 Unemployment and Recruiting Intensity Entered Together

A rise in the vacancy-to-unemployment ratio can reflect more vacancies, fewer unemployed workers, or both. Table E.3 asks whether recruiting intensity helps explain minority earnings gaps once unemployment is taken into account.

We measure recruiting intensity as the log of Lightcast postings per LAUS labor-force participant in the same MSA and quarter. Dividing by the labor force adjusts for market size without using unemployment as the denominator. The posting counts are constructed as described in Appendix C.3. We compare the preferred specification with and without recruiting intensity, using the same observations within each job population so that differences reflect the added regressor rather than changes in the sample. The exercise does not test the restriction the ratio imposes on its two components, since the unemployment side enters in percentage points and not in logs.

TABLE E.3: UNEMPLOYMENT AND RECRUITING INTENSITY ENTERED TOGETHER

	Beginning-of-quarter jobs		New stable hires	
Outcome: W - M earnings gap	Baseline	Add recruiting intensity	Baseline	Add recruiting intensity
Unemployment ($-U$)	-0.21^{***}	-0.21^{***}	-0.30^{***}	-0.29^{***}
	(0.05)	(0.05)	(0.08)	(0.08)
Recruiting intensity	—	$+0.01$	—	-0.04
		(0.12)		(0.21)

Notes: Each column reports a regression of the white-minus-minority earnings gap, constructed as 100 times the difference in log mean earnings. Entries are changes in the gap in percentage points. Baseline columns include the negative LAUS unemployment rate; the remaining columns add recruiting intensity, measured as log Lightcast postings per LAUS labor-force participant. Unemployment entries correspond to a 1-percentage-point decline in unemployment and recruiting entries to a doubling of postings per labor-force participant. All specifications include MSA-by-quarter-of-year and calendar-quarter fixed effects and MSA-specific linear trends, weight MSA–quarters equally, and cluster standard errors by MSA. The sample covers 2010Q3–2025Q2, including 2020, with 22,158 beginning-of-quarter and 21,990 new-stable-hire observations in 374 MSAs. Baseline and joint specifications use identical observations within each population, so their differences reflect the added control rather than sample changes. All-period MSA exclusions follow Table D.1. Parentheses report standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Adding recruiting intensity leaves the unemployment estimates essentially unchanged, at approximately -0.21 percentage points for beginning-of-quarter jobs and -0.30 percentage points for new stable hires per 1-percentage-point unemployment decline. The recruiting estimates are $+0.01$ and -0.04 percentage points per doubling and insignificant. The exercise therefore finds no evidence that the vacancy side adds information about the gap once unemployment is included.

E.4 Online Postings against the Official Vacancy Count

The absence of a detectable vacancy-side relationship could reflect limitations of the online-postings measure. To assess this possibility, Table E.4 repeats the comparison of unemployment and recruiting intensity using either Lightcast postings or official JOLTS openings. We conduct the exercise at the state level, where JOLTS provides coverage over our full sample period.

The state minority earnings gap is built from QWI state records using the construction method previously applied for metropolitan areas, and the unemployment measure and the labor-force denominator come from the BLS statewide LAUS series. The vacancy side enters once as Lightcast postings per labor-force participant and once as JOLTS openings per labor-force participant, each alongside the unemployment measure, on the same observations within a job population. The regressions use the state-level counterpart of our preferred specification.

TABLE E.4: ONLINE POSTINGS AND OFFICIAL VACANCY COUNTS AT THE STATE LEVEL

Outcome: W - M earnings gap	Online postings	Official openings
Population: Beginning-of-quarter jobs		
Unemployment	−0.11	−0.11
	(0.09)	(0.09)
Vacancy side	+0.18	−0.12
	(0.58)	(0.84)
Population: New stable hires		
Unemployment	−0.31	−0.31
	(0.19)	(0.19)
Vacancy side	+0.08	−0.02
	(0.42)	(0.66)

Notes: Each column reports one regression of the state minority earnings gap on the unemployment measure and a vacancy side entered together, following the metropolitan specification of Table E.3. The outcome is the state white-minus-minority earnings gap, constructed as 100 times the difference in log mean QWI monthly earnings in the indicated job population, using the same construction as for metropolitan areas. Entries are changes in the gap in percentage points. The first column measures vacancies using Lightcast online postings and the second using official JOLTS openings. The comparison is conducted at the state level, the most granular geography with JOLTS coverage over the full sample period. Minorities pool non-Hispanic Black workers and Hispanic workers of any race; white workers are non-Hispanic. Both vacancy sides enter as the log of vacancies per labor-force participant, so the unemployment entry is read for a 1-percentage-point decline in the state unemployment rate and the vacancy entry for a doubling of vacancies per labor-force participant. State unemployment and the labor force come from the BLS statewide LAUS series, and postings are the Lightcast county file summed to the state. All regressions cover 2010Q3–2025Q2 in 50 states and DC, including 2020, with state and calendar-quarter fixed effects, state-specific linear trends, and state-by-quarter-of-year fixed effects. Each state–quarter receives equal weight and standard errors are clustered by state. Within a job population both columns use the same observations, 3,010 state–quarters for beginning-of-quarter jobs and 3,005 for new stable hires. Parentheses report standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The vacancy side is statistically indistinguishable from zero under both sources and for both job

populations. The unemployment estimates remain negative and broadly comparable in magnitude to the metropolitan estimates, but are also imprecisely estimated. The absence of a detectable vacancy-side relationship is therefore not specific to online postings, although the imprecision of the state estimates limits the strength of this check.

F Assessing the Role of Composition and Policy

F.1 Worker-side Composition

F.1.1 Workforce and Population Controls

We use ACS and CPS controls to assess whether changes in worker composition account for the tightening response of QWI earnings gaps among beginning-of-quarter jobs and new stable hires (Appendix C.1). For each survey and earnings population, we use the same MSA–quarters in every specification, retaining only observations with all the controls needed for the most complete adjustment. This keeps the sample fixed as controls are added, so that changes in the tightness estimate reflect the adjustment rather than changes in which observations enter the regression.

The ACS workforce controls use survey-weighted means among civilian household residents aged 16–64 who are not students and report wage-and-salary work with usable weekly earnings, following the eligibility rules in Appendix C.2. For our first set of controls, we compute annual white-minus-pooled-minority differences in mean age; female share; high-school, some-college or associate-degree, bachelor's-degree, and postgraduate-degree shares; married share; and disability share. We pool Black and Hispanic respondents using their person weights to calculate each minority mean or share. Each annual difference is assigned to all four QWI quarters in that year. The ACS comparison covers 2012–2019 and 2021–2024.

The second set of ACS controls describes immigration history and English proficiency through white-minus-minority differences in group shares. We distinguish foreign birth, noncitizenship, and limited English proficiency; arrival in the United States 0–5, 6–10, 11–20, or 21 or more years earlier; and origins in Mexico, Central America and the Caribbean, South America, Europe and Canada, Asia, Africa, and Oceania or other regions. Each share uses all eligible wage workers in the group as its denominator, not only immigrants. Thus, these controls capture both the prevalence of immigration in each workforce and differences in immigrants' backgrounds. People born in US territories or abroad to US parents are not classified as immigrants.

To capture changes in local labor supply beyond the employed workforce, the final set of ACS controls adds all-race population shares. These are the bachelor's-or-higher and foreign-born shares among civilian household residents aged 16–64, and the share aged 55 or older among those aged 16 and older. They include students and nonworkers, with retirement-age residents included in the age measure. As with the workforce controls, each annual share is assigned to all four quarters

of its year.

The CPS workforce controls use Basic Monthly person weights among civilian nonstudent wage-and-salary workers aged 16–65. Because only worker characteristics are needed, we can use respondents beyond the outgoing-rotation subsample asked about earnings. We construct white-minus-minority differences in mean age and female, high-school, some-college or associate-degree, bachelor's-or-higher, and married shares. To reduce sampling fluctuations, we pool respondents from the current and preceding quarters and calculate weighted means and shares over those six months, rather than averaging two quarterly means. This smoothing applies only to the CPS controls; QWI earnings and unemployment remain current-quarter measures. The CPS comparison covers 2010Q3–2025Q1.

The final CPS specification adds the bachelor's-or-higher share among civilian residents aged 16–64 and the share aged 55 or older among residents aged 16 and older. As in the ACS, these all-race population measures include students and nonworkers. They use the CPS civilian noninstitutional population and the same six-month pooling as the workforce controls.

Both surveys cover fewer MSAs than the QWI because their public-use geographic identifiers do not permit a comparable match to every metropolitan area. The ACS identifies residential Public Use Microdata Areas (PUMAs), which can cross MSA boundaries. We assign each PUMA to the MSA containing most of its residents. For each MSA, we calculate the percentage of its residents excluded by this assignment and the percentage of assigned residents who live outside the MSA. We retain the MSA only if the sum of these two percentages is below 15 percent under both the 2010 and 2020 PUMA definitions. Applying the threshold to both definitions limits the risk that the PUMA redesign appears as a change in worker composition. The geographic match supports 240 MSAs; the paper-wide exclusions reduce this to 237.

For the CPS, metropolitan identifiers are available for only part of the sample, and their boundaries change over time. We match the CPS areas to the 2023 MSA definitions used in the QWI, requiring their shared territory to contain at least 90 percent of the population of each area. A match must therefore cover nearly all residents of both the CPS area and the QWI MSA, not merely place one smaller area inside a much larger one. We also require all six months used to construct a control to represent the same geographic area, so that boundary changes do not create artificial changes in composition. Areas and periods that cannot be matched on this basis are excluded. Together with data availability and the paper-wide exclusions, these requirements leave 260 MSAs

in the CPS comparison. Table III reports the resulting observation counts for each earnings population.

F.1.2 Cross-MSA Flows and Earnings Accounting

The cross-MSA accounting asks whether the number and previous earnings of movers can explain changes in local minority earnings gaps, without treating pay gains at the new job as a change in worker composition. We use release R2026Q2 of the J2J origin–destination data for not-seasonally-adjusted, all-industry employer changes between distinct MSAs, retaining moves between stable jobs with published positive counts and usable earnings at both ends. Within-MSA moves and moves involving nonmetropolitan areas are excluded. For a transition in quarter t, origin earnings refer to $t-1$ and destination earnings to $t+1$. Locations refer to employers rather than necessarily to workers' residences. We pool Black and Hispanic counts and earnings totals before calculating minority means or rates.

We combine these flows with QWI stable-job counts and earnings to compare movers with the local workforce. For a group (either whites or minorities) in a given MSA, let N^- and P^- denote its stable-job count and total quarterly earnings in $t-1$, and N^+ and P^+ the corresponding quantities in $t+1$. Because QWI mean earnings are reported per month, we multiply them by three and by the matching stable-job count to obtain quarterly earnings totals. Mean earnings before and after the move are $\mu^- = P^-/N^-$ and $\mu^+ = P^+/N^+$. Let D be the number of departures and P_D^- their pre-move earnings total; let A be the number of arrivals and P_A^- and P_A^+ their origin- and destination-job earnings totals. We omit MSA, group, and transition-quarter subscripts to simplify the notation below.

Panel A of Table IV expresses arrivals as a share of the workforce they join, $100A/N^+$, and departures as a share of the workforce they leave, $100D/N^-$. Panel B asks how movers' previous earnings compare with those of their own group in the relevant local labor market before the move. For arrivals, we compare their mean origin-job earnings, P_A^-/A, with the destination MSA's μ^-. For departures, we compare P_D^-/D with the origin MSA's μ^-. Each outcome is 100 times the log of this ratio. A decline as markets tighten means that movers' prior earnings fall relative to the local benchmark, not that moving lowers their pay.

Panel C combines the size of the flows with their earnings to measure their contribution to each group's local mean. We first calculate the pre-move mean after subtracting departures and their

earnings, and the post-move mean after subtracting arrivals and their earnings:

$$\mu^{\text{no-out}} = \frac{P^- - P_D^-}{N^- - D}, \qquad \mu^{\text{no-in}} = \frac{P^+ - P_A^+}{N^+ - A}. \tag{3}$$

We then add arrivals back to the post-move workforce at their previous pay. This replaces their destination earnings total with their origin earnings total while retaining their actual number:

$$\mu^{\text{pre-pay}} = \frac{P^+ - P_A^+ + P_A^-}{N^+}. \tag{4}$$

The departure contribution measures the change in the pre-move mean when leavers are removed. The arrival contribution measures the change in the post-move mean when arrivals are added at previous pay. Expressed as approximate percentage changes in each group's mean earnings, these contributions and their sum are

$$c_D = 100 \log\left(\frac{\mu^{\text{no-out}}}{\mu^-}\right), \qquad c_A = 100 \log\left(\frac{\mu^{\text{pre-pay}}}{\mu^{\text{no-in}}}\right), \qquad c_F = c_D + c_A. \tag{5}$$

In the arrival comparison, workers other than arrivals have the same observed earnings in both calculations. Therefore, the difference comes only from adding arrivals at their previous pay, measuring their contribution to the local mean without counting their new job earnings gains.

We estimate each outcome separately for whites and minorities, and regress its white-minus-minority difference on tightness to compare the groups' responses directly. For Panel C, this difference gives the contribution to the minority earnings gap, measured in percentage points. All regressions use the same MSA–quarters so that the comparisons across rows and groups are not affected by changes in the sample. Both groups must have the flow counts, employment counts, and earnings needed for every comparison above. We also require workers with positive total earnings to remain after subtracting departures from the pre-move workforce and arrivals from the post-move workforce; otherwise, the remaining workers' mean earnings needed for Panel C cannot be calculated.

The sample, from 2010Q3–2025Q1, ends one quarter before the core panel because the accounting requires earnings in $t + 1$. Because the arrival, departure, and combined contributions use the same observations and specification, their estimated responses also add up: the combined response equals the sum of the arrival and departure responses.

F.2 Vacancy-side Composition: Occupational Tightness Construction

To construct the occupational-tightness index in equation (2), we harmonize ACS employment, CPS unemployment, and Lightcast recruiting into eight groups: management and business; professional and technical occupations other than health; health practitioners and technical occupations; services including health support; sales and office occupations; farming, construction, extraction, installation, maintenance, and repair; production; and transportation and material moving. ACS and Lightcast classifications use SOC major groups; CPS occupations use the harmonized OCC1990 codes. Occupations outside these groups are excluded, accounting for 1.4 percent of otherwise eligible Lightcast postings in the 212-MSA sample and 0.02 percent of eligible CPS records; all eligible ACS employment records can be assigned to a group.[28] The broad groups provide a common classification across sources while preserving differences in the occupations held by whites and minorities.

The employment shares s^r_{moy} determine how much weight each occupation receives in a group's index. We calculate them from person-weighted ACS employment among civilian household residents aged 16–64 in private-sector or state and local government wage-and-salary jobs. Because the shares describe where group members work rather than their earnings, we include students and impose no earnings or hours restriction. For minorities, we pool Black and Hispanic employment within each occupation before calculating its share of total minority employment. To avoid constructing these distributions from very small samples, we require at least 100 unweighted minority employment records across 2012–2014, leaving 212 MSAs.

To reduce sampling noise while allowing the employment shares to change, we update them annually using the three most recent preceding ACS survey years. For 148 MSAs with at least 100 minority employment records in 2007–2009, the series starts with that window and uses the harmonized 2013 metropolitan definitions. The remaining 64 MSAs use the 2023 geographic assignments and first obtain shares from 2012–2014. To retain their earlier earnings observations, we also apply those first shares to outcome years 2010–2014. Only the employment weights are carried backward; recruiting and unemployment still reflect the quarter being studied. Thus, these early weights use ACS years that postdate the earnings observations they are applied to. When updating shares, we skip the unavailable standard 2020 ACS release and use the three most recent

[28] These percentages use unweighted records after the other source-specific restrictions described here. Lightcast and CPS cover 2010Q3–2025Q2; the ACS employment inputs cover 2007–2024, excluding 2020.

available survey years; this does not remove 2020 earnings observations.

To construct recruiting intensity, $v_{mot} = V_{mot} / \bar{E}_{mo}$, we assign Lightcast postings to MSAs through their county and average the monthly counts of newly dated postings within each quarter to obtain V_{mot}. The denominator $\bar{E}_{mo}$ is the occupation's average all-race, person-weighted ACS employment in 2012–2014, using the employment sample defined above. Keeping this denominator fixed ensures that recruiting intensity rises only when postings rise, rather than when the number of employed workers falls.

The occupational unemployment measure is $\widehat{u}_{mot} = U_{mt} q_{ot}$, where U_{mt} is the local LAUS unemployment rate and $q_{ot} = u_{ot}^{\text{CPS}} / u_{t}^{\text{CPS}}$ compares the national unemployment rate in occupation o with the rate across all classifiable occupations. To construct both national rates, we use CPS civilian respondents aged 16–64 who are employed or unemployed with previous work experience, retaining those with a classifiable occupation and a positive composite weight. Unemployed respondents are assigned by their previous occupation, so new entrants without one are excluded. We pool the monthly observations within each quarter and use the composite weights to calculate unemployment as a share of the corresponding labor force.

Because the same local rate U_{mt} enters both groups' unemployment indices, it cancels from R_{mt}. A change in overall MSA unemployment alone therefore does not move the relative index: the groups differ through their occupational employment shares and the recruiting and unemployment conditions of those occupations. This construction measures relative occupational conditions without requiring directly observed local unemployment separately for minorities and whites.

F.3 Policy Levers

F.3.1 Minimum-Wage Construction and Comparison

We construct quarterly state wage floors from daily rates in Vaghul and Zipperer (2022) through December 2022, extended through June 2025 using recorded state policy changes and their effective dates. For each day, we use the higher of the state and federal floors. Averaging these rates within the quarter gives each wage floor a weight proportional to the number of days it applies, rather than treating a mid-quarter increase as if it applied throughout. We then express the quarterly average in 2024 dollars by multiplying it by the ratio of the 2024 annual average CPI-U to the current quarter's average CPI-U (US Bureau of Labor Statistics 2026a).

For MSAs spanning more than one state, we construct fixed employment weights from QWI counts with valid disclosure status. Each state's weight is its share of the MSA's private-sector jobs in 2010Q2, just before the sample begins. The measure records statutory floors, not the share of workers bound by them, and does not comprehensively capture local ordinances, exemptions, eligibility differences, or compliance.

F.3.2 Diversity, Equity, and Inclusion

F.3.2.1 DEI Classification and Review

The posting-based measures distinguish employers' stated commitments and actions from their recruitment for a dedicated DEI function. The first measure, substantive DEI content, equals one when the advertisement body contains a qualifying commitment or action related to workplace diversity or inclusion. The second, specialized DEI-role advertising, identifies the occupational function of the advertised position from its title. Neither measure requires the other to be positive.

Classification rules. The algorithm applies a fixed set of phrase-matching rules. We first standardize text formatting, then look for specified terms and phrases, sometimes requiring them to appear close together and in a particular order. For example, an invitation rule looks for language encouraging applications from underrepresented groups. These restrictions make the classification more selective than a count of isolated DEI-related words, but do not amount to a contextual reading of the full advertisement.

Four groups of body rules identify actions: invitations to underrepresented applicants; outreach and recruiting pipelines; inclusive selection practices; and organizational resources, such as employee resource groups, DEI teams, or targeted mentoring. Commitment rules identify affirmative language about organizational priorities. Examples include *committed to* in close proximity to diversity or inclusion language, and statements about building or maintaining an inclusive workplace. A concrete initiative is not required for a commitment to qualify. Substantive DEI content is the union of these action and commitment indicators, so several matches in the same advertisement still count as one positive posting.

Legal boilerplate is recorded separately through equal-opportunity employer statements, nondiscrimination language, protected-characteristic lists, federal-contractor statements, and legally framed accommodation notices. These matches alone do not produce a positive substantive DEI content

classification. The substantive measure encompasses inclusion beyond race; a racial term is not required. These choices make the measure one of advertised organizational engagement, not a count of concrete initiatives or evidence that an initiative was implemented.

We search both title fields supplied by Lightcast and combine the inclusion-related and occupational terms they contain when classifying the posting; an exclusion term in either field rules it out. A title qualifies through an explicit DEI phrase or acronym; inclusion-related language combined with an occupational term such as manager or specialist; a diversity-recruiting phrase; or a designated equal-opportunity or affirmative-action role. Specified terms exclude supplier and procurement, ecological, student and academic, health, and financial or other non-DEI uses of equity. Because these are textual rules, ambiguous acronyms can generate false positives and mixed-function titles can be excluded. The measure counts advertisements for the identified roles, not all advertisements posted by employers that recruit for them.

The full posting sample contains 391.9 million advertisements across 374 MSAs over 2010Q3–2025Q2. The algorithm identifies 119,139 advertisements for specialized DEI roles and 25.75 million with substantive DEI content, representing 0.03 and 6.57 percent of the respective eligible postings.

Independent reading of advertisements. The review assesses whether the automated labels correspond to readings of the advertisements under the intended definitions. The reviewed postings come from a stratified review roster spanning 2010–2025. The roster deliberately overrepresents algorithm-positive postings and negative postings with potentially relevant language, alongside other negative cases, rather than reproducing their frequencies in the posting population. This design provides more examples of rare classifications, particularly specialized DEI roles, and potential omissions. The review is therefore informative about classification boundaries, but, because it overrepresents potentially difficult cases, its unweighted disagreement rates should greatly exceed those in the full posting population.

Three research assistants review 600 postings selected to cover different periods and combinations of algorithmic classifications. Selection balances coverage of classification patterns within periods rather than mirroring their frequency, further increasing the representation of rare cases. They receive the full available title fields and body and the coding instructions, without the algorithm's labels or other information, and provide five separate judgments: whether the posting

contains specialized DEI roles, boilerplate, substantive action, substantive commitment, and explicitly racial or ethnic substantive content. They distinguish an action from an aspiration, while allowing an explicit workforce-inclusion commitment to qualify without a concrete initiative. For specialized roles, the title is primary and the body may clarify an acronym, but cannot turn a generic title into a DEI occupation. They are instructed to read the material in context rather than implement another keyword rule and to leave unsupported judgments unresolved. Two hundred different postings are reviewed by each assistant, so the sample does not allow us to measure agreement between reviewers. Their rationales were not collected, so explanations of their disagreements are not available. Their judgments are retained without adjudication and are not treated as an error-free benchmark.

For the algorithm and the reviewers alike, substantive content is positive if either action or commitment is positive, negative only if both are negative, and unresolved otherwise. Unresolved judgments are excluded from the comparison. The comparison uses 582 resolved postings for substantive content and 593 for specialized roles.

The algorithm classifies 47.2 percent of the reviewed postings as specialized DEI roles and 66.3 percent as containing substantive DEI content, compared with 35.8 and 47.6 percent for the human reviewers (Panel A of Table F.2). These percentages describe the selected review sample, which deliberately includes DEI-positive and edge cases, rather than represent percentages for the full postings population.

Agreement is higher for specialized roles, but the direction of disagreement matters for both measures. Table F.1 reports agreement rates and the direction of disagreement for the reviewed postings.

Algorithm–human agreement is 68.9 percent for substantive content and 87.9 percent for specialized roles. The algorithm identifies more substantive content than the human reviewers. This comparison involves a more subjective judgment than identifying a specialized role, particularly at the boundary between an explicit inclusion commitment and generic employer language.

Although disagreement is not trivial, the review sample oversamples algorithm-positive postings and potential boundary cases. Panel B of Table F.2 reweights postings using their period and original sampling group, distinguishing positive cases, potentially difficult negatives, and ordinary negatives, while matching the full-sample algorithm-positive shares. The same weights apply to both. In the reviewed sample, the weighted shares are 0.03 percent for specialized DEI roles and

6.6 percent for substantive content under the algorithm, compared with 0.71 and 15.3 percent for human reviewers, respectively. Differences remain, particularly for substantive content, likely the result of the limited sample size and of the weights adjusting the composition of the review sample but not eliminating selection within the weighting groups.

TABLE F.1: VALIDATION OF POSTING-BASED DEI CLASSIFICATIONS IN SAMPLES OVERSAMPLING BOUNDARY CASES

Measure	*N*	Agreement (%)	Both positive	Algorithm only positive	Human only positive	Both negative
Substantive DEI content	582	68.90	241	145	36	160
Specialized DEI roles	593	87.86	210	70	2	311

Notes: The comparison uses the algorithmic classifications and independent human judgments without adjudication. The review sample deliberately oversamples algorithm-positive postings and algorithm-negative postings with potentially relevant language, giving greater representation to rare classifications and boundary cases. Reviewers read the available posting text and titles without the algorithm's labels. Substantive content is the union of substantive action and commitment; specialized roles concern the occupational function. Unresolved judgments are excluded separately by measure, not recoded as zero. The table draws on 600 postings reviewed by three research assistants and uses the 582 resolved substantive-content and 593 resolved specialized-role comparisons. Each posting has one human review, so agreement between human reviewers cannot be calculated. Agreement is the percentage receiving the same binary classification. The four count columns partition the resolved sample: both positive, algorithm only positive, human only positive, and both negative. Human judgments are not assumed error-free. These unweighted statistics describe the reviewed advertisements, not population prevalence or error rates. Appendix F.3.2.1 provides the definitions and review procedures; Table F.2 reports positive-classification shares.

TABLE F.2: DEI CLASSIFICATION SHARES IN SAMPLES OVERSAMPLING BOUNDARY CASES

A. Positive classifications in the review sample (%)

Measure	N	Algorithm	Human
Specialized DEI roles	593	47.22	35.75
Substantive DEI content	582	66.32	47.59

B. Positive classifications calibrated by sampling stratum (%)

Measure	N	Algorithm	Human
Specialized DEI roles	593	0.03	0.71
Substantive DEI content	582	6.57	15.27

Notes: Panel A reports unweighted positive-classification shares; Panel B reports shares calibrated by sampling stratum. Both use the resolved samples of Table F.1, drawn from 600 postings reviewed by three research assistants. The samples deliberately oversample algorithm-positive postings and algorithm-negative postings with potentially relevant language, giving greater representation to rare classifications and boundary cases. Each posting has one human review. Substantive content is the union of substantive action and commitment; specialized roles concern the occupational function. Unresolved judgments are excluded separately by measure, and each row uses the same observations for the algorithm and the reviewers. For each measure, Panel B uses 25 cells: five periods (2010–2011, 2012–2014, 2015–2016, 2017–2019, and 2020–2025) crossed with five original sampling groups (positives with one or multiple qualifying rule matches, negatives with near-target signals or other candidate-rule signals, and ordinary negatives outside the candidate set). Within algorithm-positive and algorithm-negative postings separately, weights reproduce the cell proportions recorded in the original text-sampling frame. They also match the full-study algorithm-positive shares of 0.03 percent for specialized roles and 6.57 percent for substantive content, using unrounded targets. Specifically, a posting in cell h receives weight proportional to $p_a N_h / (N_a n_h)$, where a is its algorithmic status, p_a its target share, N_h / N_a the cell's share of that status in the sampling frame, and n_h its resolved review count. Weights are positive and identical across classifiers within each row; N remains the actual number reviewed. The text-sampling frame is broader than the MSA estimation sample, and calibration does not correct residual selection within cells. These are descriptive calibrated comparisons, not population-prevalence estimates for the algorithm or the human reviewers. Human judgments are not assumed error-free. Appendix F.3.2.1 provides the review procedures and definitions.

F.3.2.2 External Validation of DEI Measures

We use three external comparisons to assess whether the DEI information extracted from the postings aligns with employer equality and inclusion engagement documented externally, outside of the advertisements. Two exercises examine formal racial-equity commitments and recruiting relationships; a third examines applicants' callback outcomes. The first two speak to organizational engagement, while the third asks whether that engagement also corresponds to less unequal applicant treatment.

To compare employers' advertised DEI engagement with the external records, we construct an employer-level share for each measure over the period relevant to each benchmark. Using the existing classifications, we divide each employer's qualifying postings by all of its postings in the linked data during that period. We report substantive content as a percentage and specialized roles as advertisements per 10,000 postings. Employers must have at least 50 matched postings and an

identified industry. We assign each employer to its most frequent two-digit NAICS category in the fourth quarter of the final year covered by each comparison, combining manufacturing, retail, and transportation categories into their standard sectors, since the employer's industry enters the selection of comparison employers.

Formal commitments and recruiting relationships. The first benchmark comprises the 35 employers listed by Management Leadership for Tomorrow (2021) as committed to pursuing Black Equity at Work certification by May 2021. This program sets standards for workplace racial equity; appearing on the list records a commitment to the process, not certification already achieved. We use 2021 posting shares and retain 26 employers with identified matches and adequate comparison support. The second benchmark is the December 2024 list of career-services sponsors and partners at North Carolina A&T, a historically Black university (North Carolina Agricultural and Technical State University 2024). Using 2024 posting shares, we retain 37 of the 40 listed employers.

For each identified employer, we select up to five employers outside the corresponding list in the same industry and closest in log posting volume. A comparison employer must have between half and twice the identified employer's posting volume, and each matched set must contain at least three comparisons. Matching is without replacement, producing 130 comparison employers for Black Equity at Work and 185 for North Carolina A&T. Employer identities and comparison groups are established without using the DEI shares, and ambiguous employer or subsidiary matches are excluded. We first subtract each set's mean comparison share from the identified employer's share, then average these differences with equal weight across identified employers. The 95 percent confidence intervals use 9,999 bootstrap resamples of matched sets.

Substantive DEI content is more prevalent among employers with independently documented inclusion commitments, most clearly in the certification-program comparison (Table F.3, Panel A). Its average posting share is 34.9 percent among those employers and 5.6 percent among their comparisons, a difference of 29.3 percentage points with a 95 percent interval of [15.2, 44.3]. The North Carolina A&T difference is also positive, at 11.2 points, but less precisely estimated: its interval is $[-1.1, 23.9]$ (Table F.3, Panel B). Specialized DEI roles are also more prevalent among employers committed to pursuing certification, by an estimated 29.0 advertisements per 10,000 postings. The North Carolina A&T comparison shows no clear difference in specialized-role advertising. Boil-

erplate is also more prevalent among the externally identified employers, so the associations are not unique to substantive language. The evidence is most informative about employers' broader commitment to inclusion rather than the implementation of a particular hiring practice.

Applicant treatment. The third exercise asks whether employers advertising more DEI engagement have smaller Black–white callback gaps. We use the difference between the callback rates of white and Black applicants reported by Kline et al. (2024). We link these outcomes to employers' DEI posting shares during the experiment's October 2019–April 2021 fieldwork window. Of the 97 employers in the external source, 92 remain after identity and posting-coverage restrictions and the exclusion of industries represented by only one employer. Each regression relates the white-minus-Black callback gap to one DEI measure, controlling for industry indicators and log posting volume and weighting employers equally. We standardize the DEI measure by its employer-level sample standard deviation, so the coefficient gives the percentage-point difference in the callback gap associated with a one-standard-deviation increase in advertised engagement.

Neither DEI measure clearly predicts smaller Black–white callback disparities (Table F.3, Panel C). The estimated response to a one-standard-deviation increase is 0.02 percentage point for substantive content and -0.3 for specialized roles, with respective 95 percent intervals of $[-0.4, 0.6]$ and $[-0.8, 0.4]$. These intervals use 3,999 employer-level bootstrap resamples. The small employer sample, however, leaves considerable uncertainty, so these estimates cannot rule out a relationship.

Taken together, the external validations support an interpretation of substantive content as advertised engagement associated with broader commitments, without establishing effective nondiscriminatory hiring.

TABLE F.3: DEI RECRUITING MEASURES AND INDEPENDENTLY DOCUMENTED EMPLOYER ENGAGEMENT

Measure	Listed employers	Comparison employers	Difference	95% CI
A. Commitment to Black Equity at Work certification				
Substantive DEI content (%)	34.87	5.56	29.31	[15.23, 44.27]
Specialized DEI roles (per 10,000)	39.32	10.35	28.97	[-0.03, 68.02]
Legal boilerplate (%)	77.65	57.91	19.74	[8.51, 31.04]
B. North Carolina A&T recruiting relationships				
Substantive DEI content (%)	26.96	15.75	11.21	[-1.10, 23.92]
Specialized DEI roles (per 10,000)	2.19	2.77	-0.57	[-2.73, 1.62]
Legal boilerplate (%)	76.12	64.43	11.69	[-0.21, 22.50]

C. white–Black callback gaps		
Measure	Response per SD (pp)	95% CI
Substantive DEI content	0.02	[-0.44, 0.64]
Specialized DEI roles	-0.31	[-0.82, 0.36]

Notes: Panels A and B compare employer posting shares with matched employers outside the corresponding published list. Panel A uses 2021 shares for 26 employers and 130 comparisons; the external cohort records commitment to pursue certification, not certification achieved (Management Leadership for Tomorrow 2021). Panel B uses 2024 shares for 37 employers and 185 comparisons from the university's December 2024 sponsor and partner list (North Carolina Agricultural and Technical State University 2024). Matches share a broad two-digit industry and have similar posting volumes, within a factor of two. Each listed employer is matched to up to five comparison employers, without replacement, and must have at least three. Listed-employer means, matched-set comparison means, and their differences weight listed employers equally. The 95 percent confidence intervals resample matched sets 9,999 times. For substantive content and legal boilerplate, means are percentages and differences and their confidence intervals are in percentage points; specialized roles are advertisements per 10,000 postings. All three use matched postings as their denominator, rather than the measure-specific nonmissing-text counts used in the MSA regressions. Panel C uses 92 employers from Kline et al. (2024), with posting shares during October 2019–April 2021. Separate regressions relate published raw white-minus-Black callback gaps to each DEI measure, controlling for broad industry and log posting volume and weighting employers equally. Entries represent one employer-level standard deviation more DEI exposure; the outcome is in percentage points. The 95 percent confidence intervals use 3,999 employer-level bootstrap resamples, holding the exposure standard deviation at its original-sample value. The source's callback rates are rounded to 0.1 percentage point; the partly industry-adjusted grades are not the outcome. Appendix F.3.2.2 describes coverage and interpretation.

F.3.2.3 Local DEI Exposure and Earnings-Gap Regressions

The earnings regressions use DEI shares to describe engagement in the local recruiting environment. Advertisements, excluding those for fully remote positions, are assigned to MSAs by posting location. For measure d, let N_{mt}^{d} be the number of advertisements flagged with either substantive DEI content or hiring for specialized DEI jobs in MSA m and quarter t, and $N_{mt}^{\text{eligible},d}$ the corresponding eligible count of total advertisements. The trailing-four-quarter share is:

$$D_{mt}^{d} = \frac{\sum_{j=0}^{3} N_{m,t-j}^{d}}{\sum_{j=0}^{3} N_{m,t-j}^{\text{eligible},d}}. \tag{6}$$

Pooling the advertisement counts smooths fluctuations in uncommon DEI recruiting while giving each eligible advertisement the same weight within the four-quarter window. Only DEI exposure is smoothed: tightness and earnings refer to the current quarter.

Across all MSA–quarters with complete four-quarter shares in 2011Q2–2025Q2, specialized-role advertising averages 2.0 advertisements per 10,000 postings, with a standard deviation of 2.6 across MSA-quarters, while substantive content averages 5.1 percent, with a standard deviation of 3.8 percentage points. The 25th and 75th percentiles are 0 and 3.0 advertisements per 10,000 for specialized roles, and 2.5 and 6.7 percent for substantive content. These local summary statistics weight MSA–quarters equally and describe the constructed measures before the outcome-specific regression restrictions. To make coefficients comparable across the two measures while preserving the time-varying share nature of the original variable, we center each share at the median and divide it by the interquartile range of pooled 2017–2019 shares across 374 equally weighted MSAs.

G Transition Frequency and Earnings-Gap Compression

G.1 Transition Rates from Job-to-Job Flows

Panel A of Table VIII uses the published, not-seasonally-adjusted metropolitan rates from the LEHD Job-to-Job Flows, which record transitions involving a worker's main job. Each rate divides the quarterly flow by average main-job employment at the beginning and end of the quarter, and therefore describes transitions per job held. Averaging the two endpoints follows the Job-to-Job Flows convention: employment changes within the quarter, and the average sets the base each flow is measured against. Direct employer changes involve no observed nonemployment spell. Hires from nonemployment and separations into nonemployment are the full nonemployment-flow measures, not the narrower persistent nonemployment measures. Nonemployment here is defined from employment records and does not distinguish unemployment from being outside the labor force. Hires are assigned to the MSA of the destination employer and separations to the MSA of the origin employer. Moves across MSA boundaries can therefore affect local hiring and separation rates differently.

The pooled minority rate averages Black and Hispanic rates using their QWI beginning-of-quarter employment as weights, which is a pooling of published rates and not a reconstruction from group-specific flow counts. The third column of Panel A is the white rate minus that pooled rate, a difference we form before estimation and estimate in its own regression. Each rate enters

equation (1) as the dependent variable. The white columns are estimated on observations with positive white employment and the minority and difference columns on observations with positive minority employment.

G.2 Initial and Stable Hiring

The final two rows of Panel A use the QWI rather than the Job-to-Job Flows and are not restricted to workers' main jobs. Initial hires are current-quarter new hires, excluding recalls to the same employer within the preceding year, relative to beginning-of-quarter employment: $100HirN_t/Emp_t$. Subsequent stable hiring is $100HirNS_{t+1}/Emp_t$, counting entry-quarter hires whose jobs reach full-quarter status in the following quarter, as defined in Appendix C.1. Both rates therefore describe the same cohort of entrants at two stages, using employment in the entry quarter as their denominator.

Black and Hispanic hires and employment are summed before the ratio is taken. The comparison uses the common availability sample for the new-stable-hire earnings regressions with flow controls, further requiring initial-hiring counts, positive entry-quarter employment, and following-quarter stable-hiring counts for both groups. Both hiring rows and all three reported outcomes use the same MSA–quarters, but not the full sample used for the J2J transition rates. Because stable status is recorded one quarter after entry, this sample ends in 2025Q1. A smaller response of stable hiring does not by itself establish a decline in retention, since retention depends on the proportional changes in stable and initial hiring, not on the difference between their percentage-point responses.

G.3 Flow Controls in the Earnings Regression

Table IX returns to the earnings outcome and adds the transition rates as controls. Every control is a white-minus-minority difference in a current-quarter rate, measured in percentage points. The employer-change control is the direct employer-change separation rate, and the nonemployment controls are the hiring-from-nonemployment and separation-to-nonemployment rates, all three from the Job-to-Job Flows and pooled as described above. The stable-hiring control is the QWI new-stable-hire rate relative to beginning-of-quarter employment, $100HirNS_t/Emp_t$. Unlike the forward-aligned measure in Appendix G.2, it records stable hiring in the earnings quarter rather than following that quarter's entrants into the next quarter.

Every column of the table, including the baseline, uses the same MSA–quarters within each earnings population, retaining only observations with all four flow controls available. Differences across columns therefore reflect the controls rather than changes in the estimation sample. Because transition rates and the composition of movers may themselves respond to tightness, these comparisons show how conditioning on measured flow differences changes the tightness estimate; they do not identify a causal share of earnings compression attributable to mobility.

H Transition Quality and Earnings-Gap Compression

The exercises in Section 8 distinguish changes in hiring destinations from upgrading relative to the jobs workers leave. Table X uses QWI initial hires to describe where hiring occurs, while Table XI uses J2J origin–destination data to account for both endpoints of employer changes.

H.1 Destination Pay Ranks

Panel A places industries, firm-size classes, and counties within each MSA on separate pay rankings. For each dimension, we average destination ranks using initial-hiring counts, without conditioning on a worker's previous employer or employment status. All ranks use a zero-to-100 scale and the same pay ranking for white and minority hires. The outcomes therefore describe the relative pay positions of each group's hiring destinations, rather than the improvement relative to the jobs individual workers leave.

All destination pay ranks use all-worker private-sector stable-job earnings over the three preceding calendar years, deflated using the CPI-U and weighted by stable-job counts. We update these ranks annually. A pay benchmark requires all three reference years, at least eight valid quarters in total, and at least two in each year. Changes in average destination rank can therefore reflect both shifts in where hiring occurs and changes in the relative pay rankings of those destinations.

Industries are ranked nationally, while counties are ranked within each MSA to capture differences in pay among locations in the same local labor market. A geographic comparison requires at least two ranked counties. Firm-size classes are also ranked within each MSA. The five classes contain 0–19, 20–49, 50–249, 250–499, and 500 or more employees; they are ranked by earnings, not assumed to pay more as firm size increases. The joint job-type ranks of Panel B use the same rolling earnings benchmarks and support requirements.

Because a destination can rank highly in one dimension but poorly in another, Panel B ranks

36 job types defined jointly by industry and firm size. Eight industries are split into three size groups—fewer than 50, 50–249, and 250 or more employees—while the remaining twelve industries enter without size splits.[29] The first row ranks these job types across metropolitan areas, where a destination's position reflects the type of job together with its location. The second ranks them within each metropolitan area to detect shifts toward better-paying local jobs that differences in pay across locations may obscure.

Black and Hispanic hire counts and rank numerators are pooled before the minority mean is formed, and the third column estimates the white-minus-minority difference in its own regression. Within each measure, both groups are compared over the same categories with released hiring counts and usable pay ranks. These categories must cover at least 75 percent of each group's total initial hires in the MSA–quarter, so that the comparison does not rest on a small observed subset of hiring.

H.2 Separation between Whites and Minorities in Hiring

Panel C measures how differently white and minority hires are distributed across the same 36 job types, rather than comparing their average positions on a pay ranking. The measure is 100 times one half of the summed absolute differences between the two groups' hiring shares, which equals zero when the two distributions coincide and rises as they diverge. A negative response therefore records convergence between them, but leaves the level of pay at those destinations open, since two groups may share a distribution anywhere on the ranking. It also does not identify which group's distribution moves toward the other's.

Hiring shares are calculated within each MSA–quarter over job types with released counts for both groups, again requiring these types to cover at least 75 percent of each group's initial hires. Unlike the rank measures, the dissimilarity index does not require a pay benchmark. It therefore retains a larger sample than Panel B: 317 metropolitan areas instead of 300 areas.

H.3 Earnings Gains across Transitions

Table XI extends the destination comparisons by accounting for the pay of the jobs movers leave. We ask whether tightening produces larger gains for minority than for white movers through

[29] The eight split industries are construction, manufacturing, wholesale trade, retail trade, professional and technical services, administrative and support services, health care and social assistance, and accommodation and food services.

changes in their origin and destination categories, valued using recent category pay. The Job-to-Job Origin-Destination statistics report mover counts and mean earnings before and after employer changes. We combine two routes: direct changes between stable jobs and stable hires following brief nonemployment. Earnings endpoints are the quarter before and the quarter after the transition for direct changes, and two quarters before against one quarter after for the brief-nonemployment route, which allows for the intervening spell. In both routes, earnings refer to a full quarter in each job; stable refers to employment continuity, not full-time work.

Both columns use cells that jointly identify origin and destination industries and MSAs. We retain a route within a cell only when its mover count and earnings at both endpoints are released and positive, and rolling industry and metropolitan pay benchmarks are available at both endpoints. The same retained cells and mover counts therefore underlie both decompositions. We apply the same transition-year benchmark to both endpoints of each move. Moves within an industry or MSA are included as well as moves between them; the corresponding category-pay component is zero for a move within the category.

Within each origin MSA, quarter, and group, we pool mover counts and earnings totals across both routes before taking logs. Let n_j be the count in retained cell–route j, and let e_j^O and e_j^D denote its mean origin and destination earnings. Suppressing the MSA, quarter, and group subscripts, the observed gain is

$$G = \ln\left(\frac{\sum_j n_j e_j^D}{\sum_j n_j}\right) - \ln\left(\frac{\sum_j n_j e_j^O}{\sum_j n_j}\right). \tag{7}$$

Black and Hispanic counts and earnings totals are likewise combined before calculating the minority gain. Thus, G is the log change in the retained movers' mean earnings, not an average of cell-level log changes or of individual workers' earnings growth. Pooling first avoids making the observed gain depend mechanically on whether the same movers are grouped by industry or geography.

The first row reports the tightening response of the white-minus-minority difference in G, with estimates and standard errors multiplied by 100 to express the response in percentage points. For the second row, we apply the same earnings-gain calculation and mover weights after replacing e_j^O and e_j^D with the reference pay of the corresponding origin and destination categories. Reference pay uses all-worker private-sector stable-job earnings over the three calendar years preceding the transition, deflated using the CPI-U and weighted by stable-job counts. The benchmarks update

annually and follow the same support requirements as the destination ranks. Industry reference pay is calculated across MSAs, while metropolitan reference pay uses all-industry earnings within each MSA. This component values moves between categories at their recent relative pay, allowing category pay differences to evolve across transition years. The third row is the observed gain net of that component, so the first row equals the sum of the second and third. This residual is an accounting remainder, not a separate estimate of within-firm wage changes: earnings can also vary with the employers and workers involved, hours, and other features not captured by category pay.

I Monopsony Power

I.1 Measuring Recruiting Concentration

We follow the vacancy-based approach of Azar et al. (2020) to measure how recruiting opportunities are distributed across employers seeking workers in the same local occupation. We adapt their commuting-zone geography to MSAs to match our earnings data, assigning Lightcast postings using the Census Bureau's 2023 metropolitan delineations and defining occupations at the six-digit SOC level. Following their methodology, within each MSA–occupation–quarter, we calculate an HHI as the sum of squared employer posting shares. Advertisements without an identified employer do not enter these shares, and fully remote postings are excluded from the main measure.

The occupation-level HHIs are combined into an MSA concentration index, C_{mt}, ranging from zero to one. For outcome quarter t, the quarterly HHIs are first averaged within each occupation over $t-4$ through $t-1$, and then averaged across occupations using their total posting counts over those four quarters as weights. Calculating the HHI separately each quarter avoids treating employers recruiting at different times as if they competed for workers simultaneously. An occupation can contribute even if its HHI is available for fewer than four quarters.

Occupations with a calculable HHI account for almost all non-fully-remote postings over the four-quarter window. Coverage averages over 99 percent across MSA–quarters, with a minimum of about 83 percent.[30]

[30] Incomplete coverage arises when postings lack an occupation code or belong to occupations for which no employer can be identified over the four-quarter window.

I.2 Benchmarking the Measure

We replicate the design of Azar et al. (2020) to check whether our Lightcast data reproduce their concentration and pay patterns. Following their work, we use commuting zones and restrict the data to 2016. The construction of the HHI indexes follows the description in the previous section. For the purposes of the replication, we follow the treatment of unidentified employers in Azar et al. (2020), assigning each such posting to a separate employer to obtain a lower-bound HHI, rather than excluding these postings as in our main measure. Similarly, we follow their exclusion of internships and retain all remote-work categories. Our main measure instead retains internships and excludes fully remote postings.

Panel A of Table I.1 compares the resulting concentration distribution with the later public data release documented by Choi and Marinescu (2024). The original study uses the 200 occupations with the most postings, but its exact occupation list cannot be recovered from the public release. We therefore select the 200 largest occupations by posting volume in our 2016 data. Of these, 160 occupational codes appear in the public release. Each pair contributes its annual-average HHI with equal weight. Thus, we follow the original design using a newer data release and a reconstructed occupation list, rather than exactly the same postings and occupations as the original study.

Panel B compares our estimates with the original paper's published relationship between concentration and pay. We replicate the first-column specifications in its Table 2 for advertised salary and OEWS hourly pay, regressing log pay on log HHI with calendar-quarter fixed effects. These comparisons follow the same specification and pay concepts.

TABLE I.1: BENCHMARKING RECRUITING CONCENTRATION AND ITS RELATIONSHIP WITH PAY

Panel A. Recruiting concentration in 2016				
Source	Mean	p25	Median	p75
Later public Lightcast snapshot	4155.44	1092.02	3333.33	6750.00
Current Lightcast replication	4183.38	1194.80	3333.33	6731.52

Panel B. Pay and recruiting concentration		
	Published estimate	Current replication
Outcome: Log Lightcast advertised salary		
Log HHI	−0.05	−0.06***
		(0.00)
Outcome: Log OEWS hourly pay		
Log HHI	−0.07	−0.06***
		(0.00)

Notes: Panel A reports unweighted summary statistics for 2016 commuting-zone–six-digit-occupation HHIs on a zero-to-10,000 scale, averaging available quarterly indices. The later public Lightcast snapshot contains 91,909 markets and 160 occupations; the current reconstruction contains 119,223 markets and 200 occupations. Each posting with an unidentified employer is treated as a separate employer; the main analysis instead uses identified-employer HHIs. Internships are excluded and remote-work categories retained. The benchmark follows Azar et al. (2020); the public snapshot is documented by Choi and Marinescu (2024). Panel B reports elasticities from unweighted regressions of log pay on log HHI with calendar-quarter fixed effects. Published coefficients are from Azar et al. (2020). Current salary and hourly-pay regressions use 203,329 and 174,173 market–quarters, respectively. Heteroskedasticity-robust standard errors are in parentheses; published standard errors and significance symbols are not reproduced. For current estimates, $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Appendix I.2 details the construction and differences in coverage.

Our data closely reproduce the concentration distribution in the public release and the pay relationships in the original study. Mean HHI is 4,183 in our data, compared with 4,155 in the public release; both have a median of 3,333 and similar upper quartiles. Our pay elasticities are approximately -0.06 for both advertised salary and OEWS hourly pay, close to the published estimates of -0.05 and -0.07, respectively.

I.3 Specifications and Reported Responses

To assess whether concentration shapes the earnings-gap response to tightening, our linear specification adds concentration and its interaction with tightness to the baseline earnings regression. We also ask whether this difference in responses depends on how tight the market already is. For this, we extend the cubic specification described in Section 5.3, interacting concentration with each of the three tightness terms.

We estimate the relationship separately for beginning-of-quarter jobs and new stable hires. Centered tightness is $x_{mt} = -U_{mt} - \overline{-U}$, where unemployment is measured in percentage points. We center concentration at its estimation-sample median and divide it by its interquartile range,

$c_{mt} = (C_{mt} - C_{50})/(C_{75} - C_{25})$, separately for each earnings population. The cubic specification is

$$y_{mt} = \alpha_{m,q(t)} + \lambda_t + \gamma_m t + \eta c_{mt} + \sum_{k=1}^{3} (\beta_k + \delta_k c_{mt})\, x_{mt}^k + \varepsilon_{mt} \tag{8}$$

where y_{mt} is the minority earnings gap defined in equation (1), and the fixed effects follow our main design in Section 4. The specification is cubic in tightness, while concentration enters linearly, both on its own and in its interactions with tightness.

For the cubic specification, the earnings-gap response to further tightening at a given level of tightness and concentration is

$$D(x, c) = \sum_{k=1}^{3} k\,(\beta_k + \delta_k c)\, x^{k-1}. \tag{9}$$

This is the slope of the estimated earnings-gap relationship at that point, expressed per 1-percentage-point unemployment decline. It need not equal the change from an actual 1-percentage-point decline, because the slope changes as unemployment falls along the cubic. The response is measured in percentage points of the minority earnings gap; a negative value indicates compression.

We evaluate this response at fixed low- and high-concentration benchmarks. These are the 25th and 75th percentiles in the 212-MSA sample with ACS occupational information available for the occupational-tightness exercise in Section 6.2. They correspond to HHIs of 942 and 2680 on the zero-to-10,000 scale. These values determine where we evaluate the response, not which MSAs enter the regression: the concentration regressions use the larger sample of 368 MSAs with the required earnings, tightness, and concentration data. Keeping the benchmarks fixed makes low and high concentration comparable across earnings populations and market conditions.

We test differences between these concentration levels using the estimated interaction coefficients and standard errors clustered by MSA. For the cubic specification, the high-minus-low response difference at tightness x is $(c^H - c^L)(\delta_1 + 2\delta_2 x + 3\delta_3 x^2)$, where c^L and c^H express the two HHI benchmarks on the centered and scaled concentration measure defined above.