Labor Market Concentration, Earnings Inequality, and Earnings Mobility*

Kevin Rinz

Center for Economic Studies
US Census Bureau

February 7, 2019

Abstract

Using data from the Longitudinal Business Database and Form W-2, I document trends in local industrial concentration from 1976 through 2015 and estimate the effects of that concentration on earnings outcomes within and across demographic groups. Local industrial concentration has generally been declining throughout its distribution over that period, unlike national industrial concentration, which declined sharply in the early 1980s before increasing steadily to nearly its original level beginning around 1990. Estimates indicate that increased local concentration reduces earnings and increases inequality, but observed changes in concentration have been in the opposite direction, and the magnitude of these effects has been modest relative to broader trends; back-of-the-envelope calculations suggest that the 90/10 earnings ratio was about six percent lower and earnings were about one percent higher in 2015 than they would have been if local concentration were at its 1976 level. Within demographic subgroups, most experience mean earnings reductions and all experience increases in inequality. Estimates of the effects of concentration on earnings mobility are sensitive to specification.

Keywords: Monopsony, labor market concentration, industrial concentration, local labor markets, earnings, earnings inequality, earnings mobility
JEL Classification Codes: J31, J42

*This paper is released to inform interested parties of research and to encourage discussion. Any opinions and conclusions expressed herein are those of the author and do not necessarily reflect the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The statistical summaries reported in this paper have been cleared by the Census Bureau's Disclosure Review Board, release authorization numbers CBDRB-FY18-469 and CBDRB-FY18-496. I thank Zach Brown, Anne Case, Martha Gimbel, Javier Miranda, Lee Tucker, Victoria Udalova, John Voorheis, Abigail Wozniak, and seminar participants atthe Census Bureau, Colby College, the Bureau of Labor Statistics, and the Sundance Conference on Monopsony in Labor Markets for helpful comments and discussions. Contact: kevin.rinz@census.gov

1 Introduction

The idea that employers are not simply price-takers in the labor market but may have the power to set their workers' wages is old, but the possibility that monopsony power could have substantial influence on economic outcomes has received renewed attention of late.[1] This attention comes as various measures of concentration and market power at the national level increase alongside stagnant wage growth and a declining labor share of income (Autor et al., 2017; De Loecker and Eeckhout, 2017; Edmond et al., 2018; Grullon et al., 2018; Hall, 2018; Traina, 2018). Policymakers have also taken an interest in the subject, with the White House Council of Economic Advisers highlighting competition issues generally and monopsony in particular in issue briefs (2016a; 2016b).

The many sources of labor market monopsony make it a particularly interesting economic phenomenon and a potentially difficult policy issue to address. While limited competition among a small number of firms in a labor market is a canonical example of a source of monopsony power, it can also arise from frictions in the labor market that make it difficult to find or accept new employment opportunities (Manning, 2003). Employer practices such as requiring non-compete agreements with employees (U.S. Department of the Treasury, 2015; Starr et al., 2018) or establishing no-poaching agreements with competitor firms (Krueger and Ashenfelter, 2018) can create these frictions and increase employers' wage-setting power, but more general, naturally occurring frictions associated with job search, geographic mobility, or heterogeneous preferences over job characteristics can also give rise to monopsony power.

Many of these sources of monopsony power are both universal and hyper-local. They are experienced broadly across industries, occupations, and geographies because they arise from fundamental characteristics of the economy, but the precise manner in which they are

[1]Smith (1776) describes a "tacit, but constant and uniform combination" among employers to control workers' wages. Robinson (1933) formalized the case of wage-setting power arising from there being a single buyer of labor in a market, coining the term "monopsony."

experienced depends on individual workers' particular circumstances, including their locality. These facts provide good reason to believe that the effects of monopsony power might be both widespread and different across groups. Here, I will document the degree to which monopsony power is prevalent within local labor markets and estimate its effects on earnings outcomes across the earnings distribution within and across demographic groups.

Empirically, research has identified a wide variety of settings in which monopsony power may be relevant to workers' economic outcomes. These include specific labor markets, such as markets for teachers (Landon and Baird, 1971; Luizer and Thornton, 1986; Falch, 2010; Ransom and Sims, 2010), nurses (Staiger et al., 2010; Matsudaira, 2014), engineers (Fox, 2010), retail workers (Ransom and Oaxaca, 2010; Dube et al., 2018a), judicial clerks and medical residents (Priest, 2010), and professional baseball players (Humphreys and Pyun, 2015); Amazon's Mechanical Turk platform (Dube et al., 2017, 2018b); the franchise sector (Krueger and Ashenfelter, 2018); and historical settings such as turn-of-the-century coal mining (Boal, 1995) and sharecropping (Naidu, 2010). There is also growing evidence that issues of imperfect competition in the labor market are broadly applicable beyond the specific institutional settings of particular labor markets (Manning, 2011; Azar et al., 2017; Dube et al., 2017; Tucker, 2017; Azar et al., 2018). Some of this broader work has found high levels of concentration in local labor markets in recent years. This fact, in combination with increases in national measures of market power, has fueled speculation that local concentration has been increasing over time. However, research has not provided evidence on this possibility, in part because of the difficulty of obtaining suitable data for investigating it.

The broad theoretical applicability of monopsony power, in combination with its demonstrated empirical relevance and the increased salience of competition and market power issues more broadly have led some to consider it a possible contributing factor in the rise of inequality over the last few decades (Council of Economic Advisers, 2016b). The periods during which inequality and measures of market power such as markups have risen overlap significantly. While the available market power measures generally do not directly reflect monopsony power, the leap to thinking that monopsony might play a role is small, since it acts directly on workers incomes, and income mobility has at best stagnated over this period (Chetty et al., 2016).

Increases in monopsony power may increase inequality in a literal sense, by changing the shape of the earnings distribution, but they may also have effects that vary across groups of workers. Webber (2015) found that increased employer power in the labor market increases inequality in the overall earnings distribution, but did not consider heterogeneity on other dimensions such as demographic characteristics. Others have considered the effects of monopsony power on specific subgroups of workers, finding, for example, that it reduces the wages of immigrants in Germany (Hirsch and Jahn, 2015), increases the gender wage gap in both Germany (Hirsch et al., 2010) and the United States (Webber, 2016), and reduces the wages of skilled workers and trainees in Switzerland (Muehlemann et al., 2013). No previous study has combined these two types of possibly heterogeneous effects by considering distributional effects within and across demographic groups, or considered related measures of earnings mobility. Here, I combine comprehensive administrative data on firms and individuals with demographic information obtained from surveys to do just that.

In this paper, I document trends in local labor market concentration in the United States between 1976 and 2015 using the Census Bureau's Longitudinal Business Database (LBD). I define local labor markets as intersections between industries and geographies, focusing primarily on four-digit North American Industry Classification System (NAICS) industries within commuting zones. I measure concentration using the Herfindahl-Hirschman Index, constructed using employment. Then, by combining LBD data with earnings information from Form W-2 and demographic information from the Census Numident, the decennial census, and the American Community Survey (ACS), I estimate the effects of local industrial concentration on earnings outcomes for various groups of workers, as well as for the workforce as a whole.

Though trends in national measures of other forms of concentration may have contributed to recent interest in monopsony, trends in local industrial employment concentration have differed substantially from trends in national industrial employment concentration over the last four decades. While mean national industrial concentration declined sharply in the early 1980s, it began increasing rapidly again around 1990 and continued to do so until the onset of the Great Recession, nearly returning to its initial level. Local industrial concentration, on the other hand, has been declining fairly consistently since 1976, with limited interruptions. By 2015, average local concentration had declined to about three quarters of its 1976 value.

The divergence between local and national industrial concentration is not sensitive to the industrial classification scheme, geographic definition of local, level of industrial aggregation, or use of employment weights. The divergence appears to be driven by differences in within-industry concentration when it is measured nationally versus locally. In a counterfactual exercise that varies different components of the national and local trends in isolation, the actual national trends tracks closely with the trend that would have been observed if only within-industry concentration had changed since 1976. Counterfactual trends based on varying within-industry concentration and local industrial composition track the actual local trend well. Imposing changes in within-industry concentration on local markets that are proportional to those experienced in national markets produces a dramatically different local concentration trend. Suggestive evidence indicates that national and local measures of concentration may differ because large national firms have extended their reach into additional markets over time while also increasingly participating in the same local markets.

Though local concentration levels differ across regions, trends are fairly similar. At the market level, though, there is substantial variation in the magnitude of changes in concentration over time. Though the local industrial concentration distribution has consistently tightened over time, especially at high percentiles, some specific markets have experienced large increases in concentration, while others have seen large decreases.

I use the substantial variation in local industrial concentration over time to estimate its effects on earnings, inequality, and mobility. Consistent with other recent research, I find that increased concentration reduces earnings. My estimates imply that moving from the median to the 75th percentile of the employment-weighted local industrial concentration distribution would reduce earnings by about ten percent. Moving from the median to the 25th percentile would increase earnings by a similar amount. Estimates produced without weighting by employment are larger in magnitude than the baseline estimates, which indicates that earnings reductions associated with increased concentration are larger in smaller markets.

I also find that the effects of concentration vary across groups of workers. First, looking across the earnings distribution, I find that increased concentration leads to greater inequality as measured by the ratio of the 90th percentile of the earnings distribution to the 10th percentile (the 90/10 earnings ratio). By comparing changes in the 50/10 earnings ratio and the 90/50 earnings ratio, I estimate that about 60 percent of the increase in the 90/10 earnings ratio arises from changes between the median and the 10th percentile. Moreover, I estimate elasticities of particular percentiles with respect to concentration and find that lower percentiles are more negatively responsive to changes in concentration than are percentiles in the middle of the distribution. Percentiles higher in the distribution change little in response to changes in concentration. My estimates are consistent with Webber (2015), in which a similar analysis was performed using individual-level unconditional quantile regressions.

While these estimates indicate that increased concentration does indeed reduce earnings and increase inequality, combining them with the changes in concentration that have actually been observed since 1976 suggests that local labor market concentration has not been a major contributing factor to broader changes in inequality and earnings growth. According to back-of-the-envelope calculations, average annual real earnings were about 1.2 percent higher and the 90/10 earnings ratio about 6.3 percent lower in 2015 than they would have been if local concentration were at its 1976 level.

The availability of demographic information from survey and administrative data sources allows me to evaluate whether the effects of local industrial concentration vary across groups defined by individual characteristics, as well. I find that the effects of concentration on average earnings are negative across most groups defined by age, race, sex, and education. The only groups for which the earnings effect point estimate is positive are women and Black workers. Notably, these groups have historically experienced significant labor market discrimination in the United States, and changes in related behaviors could rationalize positive aggregate earnings effects for these groups.

All demographic groups experience increases in inequality when concentration increases. Men, older workers, and workers with high school diplomas or less see the largest increases in the 90/10 earnings ratio. As in the overall distribution, these increases are generally driven by the bottom of the distribution. Women and Black workers are again exceptions, with virtually all of the inequality increases in these groups coming from the top half of the distribution. This could be due in part to the fact that these groups generally have lower earnings throughout the distribution. As a result, changes experienced at any given point in the overall earnings distribution are experienced further up the distribution of earnings withing these groups.

Finally, I estimate the effects of local industrial concentration on earnings mobility over horizons extending up to five years. My baseline estimates indicate that increases in concentration reduce relative mobility and increase absolute mobility. However, these estimates are more sensitive to specification changes, which cautions against drawing strong conclusions about earnings mobility effects at this point.

This rest of this paper proceeds as follows. Section 2 discusses measurement issues and describes the data I will use to investigate these questions. Section 3 lays out trends in local industrial concentration over four decades. Section 4 describes my approach to estimating the effects of local industrial concentration on earnings and inequality. Section 5 reports results, and Section 6 discusses them and concludes.

2 Data and Measurement

Two important questions must be answered before considering trends in local labor market concentration or the effects of concentration on earnings, inequality, and mobility. First, what constitutes a local labor market? This is, of course, a question of very broad interest, and resolving it is well beyond the scope of this paper. Fundamentally, the definition should capture the set of reasonable potential employers for a given worker. Common approaches include using geographies such as county or commuting zone, job characteristics such as industry or occupation, or interactions among these to define local labor markets. Here, I use interactions between industry and geography to define local labor markets. I discuss this further below.

Second, how can we measure local labor market concentration and the outcomes of interest? Some business data are available publicly, but they do not provide firm-level information with fine geographic detail, limiting their usefulness for measuring local employment concentration. As for outcomes, few local labor markets are sufficiently well represented in surveys to construct reliable distributional statistics. Fortunately, I can address both of these issues using administrative records available through the U.S. Census Bureau. The Bureau's data linkage infrastructure also allows me to construct earnings measures that incorporate demographic information available from the Census Numident file, the 2000 and 2010 decennial censuses, and all available years of the American Community Survey. The rest of this section details the relevant datasets and how they figure into my analysis.

2.1 The Longitudinal Business Database

The Longitudinal Business Database (LBD) provides key information such as employment, payroll, location, industry, and firm affiliation on an annual basis for all employer establishments in the United States (Jarmin and Miranda, 2002). Data, which are compiled from the Business Register (BR), the Economic Census, and other surveys, are available from 1976 through 2015 and cleaned to facilitate easy linking over time, with the database containing one observation per establishment per year.

The availability of firm identifiers, in combination with employment, industry, and geography information, permits the construction of firm-based measures of employment concentration within industry-by-geography cells. As these cells are intended to approximate labor markets here, there are some conceptual questions about what the appropriate levels of aggregation are when constructing these measures. For example, what level of geographic aggregation is appropriate? Previous studies of local labor markets have used areas as small as counties and as large as states, as well as intermediate constructions such as metropolitan areas and commuting zones. Empirically determining the ideal construction of local labor markets is beyond the scope of this paper; I use commuting zones as my preferred geographic unit.[2]

The appropriate level of industrial aggregation is also an open question. In product markets, using more precise industrial classifications probably identifies more reasonable sets of close competitors, but does this also identify more reasonable sets of alternative employment opportunities for workers? Could human capital be transferable across reasonably fine industry categories to a greater degree than the goods or services produced by those industries are substitutable for each other? This is to some extent an empirical question that I leave to future work, but it is also a practical question in this setting. The more precise the industrial classification used, the fewer establishments (and by extension the fewer firms

[2]As discussed below, my results are not sensitive to this choice.

and the less employment) will figure into the analysis. In order both to capture a broader set of alternative employment opportunities and to include additional establishments that cannot be classified in the most precise terms, I use the four-digit NAICS industry codes, an intermediate level of classification, in my analysis here.

Industrial classification schemes vary a great deal between 1976 and 2015. Within that period, NAICS replaced SIC as the dominant industrial classification system in the United States, and industry codes underwent periodic updates within each of those schemes to reflect changes in economic activity. Over a period as long as the 40 years used here, those changes in economic activity can be meaningful. If industry classifications serve as a proxy for labor markets, it might make sense to use contemporaneous classifications and allow the labor market definitions to vary over time as the economy changes rather than standardizing them. Also, updates to classification systems often result in at least some existing industries seeing establishments re-classified into different or new industries, making industry codes difficult to harmonize over time using aggregate crosswalks.

Using contemporaneous classifications also has drawbacks. Actual labor markets are not redefined sharply at five year intervals like industry codes are. Practically speaking, this limits the amount of temporal variation within industries that is available for use in regression models that include industry fixed effects. Moreover, in some cases, the definitions of particular industries (and therefore the set of establishments they contain) change subtly over time even as the codes used to identify them remain the same, so even longstanding industry codes do not necessarily represent consistently defined labor markets. Since this paper includes regression analysis that relies on within-industry variation, I use standardized industry codes. Although the primary period of interest in my regression analysis is 2005 through 2015, I use standardized industry codes for the full period covered by my descriptive analysis in order to present consistent information throughout.

I standardize industry codes using a set of crosswalks developed by Fort and Klimek

(2018). Rather than generating aggregate correspondences between industry codes over time or assign establishments in industries that split by randomizing, Fort and Klimek construct their crosswalks at the establishment level. They take advantage of the longitudinal nature of the LBD to bridge the transition from SIC to NAICS, resolve ambiguous re-classifications, and generate consistent industry codes. I obtain the Fort-Klimek industry code from the most recent available year for each establishment and use it to classify that establishment in all years of its operation.

Before proceeding, it is worth explicitly stating the decisions I make about these aggregation issues in my baseline analysis. I define local labor markets as commuting zone-level, standardized, four-digit NAICS industries. As I show below, comparisons between national and local trends in industrial concentration are little changed when constructed using contemporaneous industrial classifications instead of the consistent Fort-Klimek industry codes, or when constructed using three-digit (instead of four-digit) NAICS industries, or when constructed using counties instead of commuting zones.[3]

2.2 Form W-2

Employers use Form W-2 to report their employees' earnings to the IRS. The form includes identifying information for both the employer and the employee, the amount of taxable wages paid to the employee, the amount of tax withheld, and some information about certain non-taxable compensation. The extract available through the Center for Economic Studies (CES) at the U.S. Census Bureau contains the Employer Identification Number (EIN, sometimes also called the Tax Identification number, or TIN), the (uncapped) amount of wages paid, and the amount of deferred compensation paid from each W-2 filed from 2005 through 2015.[4] The

[3]Though not reported here, my regression results are also robust to estimation based on measures constructed using three-digit NAICS industries or counties. These results are available upon request.

[4]The form reports the amount of wages that are subject to the Social Security and Medicare payroll taxes. The Social Security payroll tax is capped, for example it was levied on only the first \$117,000 of wage income in 2014. The Medicare payroll tax is uncapped, i.e. it is levied on all wage income. I use

personally identifiable information (PII) contained on each form is used to assign a unique person identifier called a Protected Identification Key (PIK) through Census Bureau's Person Identification Validation System (PVS) and is then removed from the files.[5]

Below, I analyze the response of the earnings distribution among people employed in various geography by industry labor markets to changes in industrial concentration. In order to use W-2s for this purpose, I need to assign each form to a person, a place, and an industry. I aggregate earnings to the person level by summing wage and salary earnings and deferred compensation across W-2s within PIKs. The W-2 data I have access to contains the employer's EIN, but no other information about the employer, so industrial classification is not readily available. For individuals who receive multiple W-2s, I retain the EIN associated with their highest-income W-2. I use the EIN to assign an industrial classification obtained from the LBD and data described in the next section. The W-2 data also do not contain any information about the geographic location of the recipients. I obtain person-level address information from other tax data described in the next subsection. Both industry and geographic information are assigned to W-2s through a process described in Appendix A.

2.3 Other Data

As mentioned above, an important limitation of the W-2 data is that they do not contain any information on the geographic location of the forms' recipients. They do contain the same individual identifier available on other tax forms that include geographic information. Specifically, I have access to extracts from Form 1040 and a collection of Form 1099 infor-

the uncapped measure of wages subject to the Medicare payroll tax in this analysis. Note that the extract does not include all information available on Form W-2; for example, information about employer-sponsored health insurance is not available.

[5]In general, PVS assigns PIKs based on PII like social security numbers, date of birth, place of birth, name, and address. Not all records can be assigned a PIK if the available PII is of low quality, contains contradictory information, or is missing important elements, but when social security numbers are available, as they are on Form W-2, PIKs can be assigned to virtually all records. On other forms, where address information is available, the process also assigns a location identifier called a Master Address File Identifier (MAFID). See Wagner and Layne (2014) for a more detailed description of the PVS process.

mation returns. The 1040 data are available annually beginning in 1998 and contain the address from which they were filed. The 1099 data are available annually beginning in 2003 and contain the address to which they were sent. For my purposes, I am interested in each W-2 recipient's county of residence (from which the commuting zone of residence is determined). I obtain this information from this tax forms using a prioritization scheme described in Appendix A.

Similarly, the W-2 data do not contain the industry of the employer. They do contain employer EINs, which could be used to link them to other sources of business data. The LBD, which contains a relatively limited set of consistently available variables, does not include business's EINs. During the period relevant to this analysis, however, EINs are available from the BR, another source of administrative data on businesses that is linkable to the LBD. With EINs obtained from the BR added to the LBD, the same industrial classifications available in the LBD are assignable to W-2s using a process described in Appendix A. I can therefore use the Fort-Klimek industrial classification system to consistently construct both measures of industrial concentration within the LBD and statistics summarizing local industry earnings distributions by linking to W-2s.

Finally, in order to conduct an analysis of earnings outcomes for various demographic groups, I obtain data on date of birth and gender from the 2016 Census Numident file, which is generated from the Social Security Administration's Numident file and contains one record for every person issued a Social Security number. I place people into three age categories: under 25, 25-54, and 55 and older. I also obtain data on race and Hispanic origin from the 2000 and 2010 Decennial census and from the ACS from 2005 through 2015. For the sake of ensuring sample sizes are large, I use the race and Hispanic origin variables to create three mutually exclusive categories: non-Hispanic White, non-Hispanic Black, and Hispanic.[6] I exclude other, much smaller race and ethnicity groups from my analysis. Finally,

[6] I use the most recently reported race and Hispanic origin values for individuals who appear in multiple

I obtain information on educational attainment from the ACS. I use education information only for individuals who are at least 25 years of age when they appear in the ACS data. Because education information is not collected on the Decennial short form and only about 15 percent of population is covered by the ACS over the available period, education is much more sparsely available. As a result, I use only two education categories: high school or less (low education) and some college or more (high education).

3 Trends in Industrial Concentration

Before estimating the effects of local industrial concentration on earnings, inequality, and mobility, I present descriptive information on the level of industrial concentration, trends in concentration over time, and geographic differences in concentration. While a few papers have considered trends in national industrial concentration, little is known about how local industrial concentration has varied across places and times.[7] To the extent that employment concentration affects labor market outcomes, local concentration is likely to be particularly relevant because most workers do not engage in geographically wide-ranging job searches; job seekers are much more likely to apply to vacancies closer to their homes (Marinescu and Rathelot, 2018), with only about a quarter looking outside their state of residence (Sinclair, 2014). Unless otherwise noted, these estimates are constructed using employment to weight observations, so the trends described here reflect the experience of the average worker rather than the average market.

Before turning to local concentration, Figure 1 presents the average HHI across national four-digit NAICS industries from 1976 through 2015, with industries weighted according

surveys. For example, for an individual who responded to the 2010 Decennial short form and the 2013 ACS, I use the values reported on the 2013 ACS. Individuals who report being of Hispanic origin are assigned to the Hispanic category regardless of race. Non-Hispanic individuals who report multiple races are categorized according to the first reported race.

[7]Benmelech et al. (2018) report the national average of local concentration within five-year bins, measured using the HHI, but their analysis is focused on the manufacturing sector.

to total employment. Average concentration falls sharply in the early years of this period, declining by roughly 40 percent between 1976 and 1983. It then sees little change until about 1990, at which point it begins increasing, nearly reaching its 1976 level by the onset of the Great Recession. This pattern is not sensitive to measuring concentration using the HHI. Online Appendix Figure B1 shows very similar patterns emerge when concentration is measured using the top-four or top-twenty firm employment concentration ratios.

To my knowledge, other studies have not presented estimates of the average national HHI prior to 1982. Autor et al. (2017) estimate average top-four and top-twenty firm employment shares by major industrial sector using the Economic Census beginning in 1982. As shown in their Figure 4, the sectors they consider exhibit similar upward trends in concentration that are broadly consistent with the national trend reported here over the same period. Moreover, Grullon et al. (2018) report a sharp decline in the share of total U.S. employment at firms with at least 10,000 employees that occurs at the same time as the sharpest decline in the national HHI trend I report (see panel D of their Figure 1). A sizable reduction in employment shares at very large firms could have a meaningful impact on the square of their employment shares, and thereby on the employment-based HHI.

Finally, when I estimate the national HHI trend within sectors defined by collections of two-digit NAICS industries, only the services sector exhibits an especially large decrease in concentration that aligns with the national trend.[8] Specifically, the decline in concentration within services is driven by information and cultural industries (NAICS 51), which includes telecommunications industries. Notably, AT&T, the dominant firm in that industry, entered into a consent decree with the Department of Justice in 1982 that required it to divest itself of local telephone companies (Pinheiro, 1987). The availability of an economic explanation for the observed change in employment concentration should alleviate any concerns that the trend presented above is an artifact of a data processing or estimation error.

[8]Sector-specific HHI trends are presented in Online Appendix Figures B2, B3, and B4

Figure 2 presents the trend in average local industrial concentration, again measured using the HHI, averaged across commuting zone by four-digit NAICS industry markets. Markets are weighted according to employment. Local concentration also declines over the late 1970s and early 1980s, though not as precipitously as national concentration. It also generally continues declining, though more slowly, through the 1990s and even most of the 2000s before increasing during the Great Recession. Like the national trend, this pattern is also evident in the top-four and top-twenty firm concentration ratio trends, as shown in the online appendix.[9]

The divergence between the national and local concentration trends is not sensitive to any of the major decisions about how the two series are constructed. As shown in Appendix B, the same pattern emerges if trends are calculated using contemporaneous industry classifications instead of consistent classifications based on Fort and Klimek (2018) (Figure B7), if local markets are defined using counties instead of commuting zones (Figure B8), if they are defined using three-digit NAICS industries instead of four-digit industries (Figure B9), and if markets are not weighted by employment in constructing the average (Figure B10). It is worth noting, however, that the increase in local concentration observed since the onset of the Great Recession in the employment-weighted figures is clearer and more continuous when employment weights are not used, suggesting that smaller markets are becoming more concentrated even as the average worker is largely not exposed to those increases.

Why have the national and local concentration trends diverged? This question can be addressed both mechanically (which components of national and local mean concentration are changing differentially?) and economically (why are those components changing differentially?). I address the mechanical component of this question through a series of counterfactual exercises. First, I consider how the national trend has evolved. Average national

[9]An alternative local construction of this figure based on counties is presented in Online Appendix Figure B6 and tells a similar story.

concentration at a given point in time can be written

$$\overline{HHI}_t^N = \sum_i Share_{it} \cdot HHI_{it}$$

where, for industry i at time t, HHI_{it} is the HHI and $Share_{it}$ is the share of national employment in that industry. Figure 3 plots the actual national trend in average HHI, as well as two counterfactual national trends: the one that would have been realized if only within industry HHIs varied over time (i.e. if industry shares of employment remained fixed at their 1976 shares), and the one that would have been realized if only industry employment shares (or, between-industry concentration) varied over time (i.e. if HHIs remained fixed at their 1976 levels).

The counterfactual trend that is based on varying only within-industry HHIs is very similar to the actually observed trend. Prior to 2000, changes in industrial composition are generally moving the average in the same direction as changes in concentration, but may explain a small share of the decline, suggesting that changes in within-industry concentration are primarily responsible for the evolution of the national trend.

Second, I perform a similar exercise on the local concentration trend. Since the national share of employment in a given market/commuting zone-industry ($Share_{c,i,t}$) can be written as the product of the share of national employment in that commuting zone ($CZShare_{ct}$) and the share of commuting zone employment in that industry ($CZIndShare_{cit}$), the average local HHI can be written

$$\overline{HHI}_t^L = \sum_c \sum_i CZShare_{ct} \cdot CZIndShare_{cit} \cdot HHI_{cit}$$

Figure 4 presents counterfactual trends analogous to those in Figure 3 that vary each of the three components of the local concentration trend in isolation: within market HHIs, within CZ industrial composition, and the share of national employment in each commuting zone.

The actual local concentration trend is also presented for reference.

Based on the counterfactual trends, changes in both market HHIs and commuting zone industrial composition put downward pressure on the average local HHI, with their counterfactuals moving roughly in tandem through about 2000. After that, the concentration-only counterfactual trends slightly upward, while the composition-only mean continues to decline. Changes in the distribution of employment across commuting zones have little impact on the overall trend.

The most striking difference between the national and local counterfactuals is the behavior of the concentration-only series. After initially declining in both settings, it increases sharply after 1990 in the national series while increasing later and only modestly in the local series. Apart from roughly the second half of the 1990s in the national series, changes in industrial composition generally put downward pressure on both the national and local average HHI.

To further illustrate the implications of the divergence between the behavior of national and local HHIs, I conduct a third counterfactual exercise. Figure 5 presents two counterfactual trends: the trend that would have been realized if only local HHIs had changed, with local industrial composition and commuting zone employment distributions held fixed; and the trend that would have been realized if each local industry's HHI had evolved proportionally to that industry's national HHI. As one might expect based on the previous two exercises, these two counterfactuals are starkly different, with the trend based on the evolution of national industry HHIs increasing steadily after 1990, while the trend based on the evolution of local HHIs declines initially and remains lower than its starting level, similar to the actual local HHI trend. This figure makes clear that local and national HHIs have behaved very differently, especially since 1990.

But why have national and local HHIs behaved differently? Suppose a small number of firms increasingly dominate national industries while also more directly competing with each other in the same local markets. That could be consistent with increasing national concentration alongside stable, lower local concentration. And indeed, this possibility appears to have some empirical support. Figure 6 shows the number of markets (commuting zone by four-digit NAICS industry cells) that contain at least one establishment belonging to one of the five largest firms by employment in that national industry. The reach of the largest firms has been expanding over essentially the entire time series, with the number of local markets with at least one top-five firm increased from nearly 25,000 in 1976 to nearly 45,000 in 2015. Notably, the rate of expansion accelerated during the 1990s, around the same time national HHIs began to increase sharply.

Figure 7 focuses on markets containing at least one top-five firm and reports the number of top five firms competing in these markets. In 1976, just over 60 percent of markets with at least one top-five firm contained exactly one top-five firm. By 2015, that share had fallen to just over 50 percent. Notably, the bulk of the approximately ten additional percent of markets with multiple top-five firms in 2015 had three or more top-five firms, as the share of markets with two such firms was fairly stable over this period. Also, as indicated by the previous figure, those ten percent represent substantially more markets in 2015 than in 1976. Together, Figures 6 and 7 show that the largest national firms have expanded their geographic reach over the past 40 years while also increasingly entering the same local markets. The expansion of the geographic reach of these top firms accelerated around the same time that national HHIs began to increase. These patterns provide suggestive evidence that this channel merits further investigation.

I now turn my attention to changes in the distribution of local industrial concentration. Figure 8 plots trends in key percentiles of the employment-weighted local HHI distribution. The box and whisker plots present the interquartile range (box) and interdecile range (whiskers), with the mean (circle) and median (horizontal line) also plotted.

The figure makes a few important features of the distribution immediately clear. First, the distribution has a long right tail; in every year, the value of the 75th percentile is more than twice that of the median, and the value of the 90th percentile is more than twice the value of the 75th percentile. As a result, the mean HHI is consistently well above the median. Second, the distribution has tightened over time, and this appears to have been driven by changes in the top of the distribution. The value of the 90th percentile has fallen by about a third between 1976 and 2015. The values of the 75th percentile and median have also fallen, but more modestly, while the 10th and 25th percentiles have seen little change in absolute terms over this period.

3.1 Geographic Variation

Returning my focus to mean local industrial concentration, I now consider possible geographic heterogeneity. Figure 9 maps the average HHI across industries within each commuting zone in 1976, and Figure 10 does the same for 2015. In both years, the areas that are most concentrated tend to be rural. In particular, the Great Plains region has a relatively large number of highly concentrated commuting zones in both 1976 and 2015. The least concentrated markets tend to be in urban areas.

Figures 11 through 14 show how the average concentration within each commuting zone has changed over time, mapping differences in logged HHIs between select years. As Figure 11 shows, the middle of the country, from Texas and New Mexico up to North Dakota and Montana is home to some of the commuting zones where markets were becoming more concentrated at the fastest rates between 1976 and 1990, even as the national average local HHI was falling during that period. This continued to be the case between 1990 and 2005 (Figure 12), though a larger number of commuting zones outside this region also became more concentrated, including several in Florida, Appalachia, and the Pacific Northwest. Between 2005 and 2015 (Figure 13), increases in concentration were more widespread, though the magnitude of these increases was generally small in percent terms. Consistent with the national trend, the larger declines in concentration during the earlier years lead to net decreases in concentration on average in most commuting zones over the full period considered (Figure 14). Just over half of markets saw decreases in concentration between 1976 and 2015, while just over 40 percent saw increases.

To summarize my findings regarding trends in local industrial concentration, I find that local industrial concentration has generally been declining since 1976, with a few brief periods of increasing concentration, including one surrounding the Great Recession. National industrial concentration, on the other hand, initially declined before beginning to increase sharply again around 1990. The divergence between the local and national trends in industrial concentration appears to be driven by differential trends in within-industry concentration when it is measured locally versus nationally. Declining values of high percentiles have lead to a tightening of the local industrial concentration distribution over time. While at least some commuting zones in all regions experienced declines in concentration since 1976, the Great Plains region is home to many of the commuting zones with the highest levels of average concentration across markets as well as those with the largest percent increases in concentration.

I now turn my attention to the effects of local industrial concentration on earnings, inequality, and mobility.

4 Estimation

As illustrated in the previous section, there is a great deal of variation in industrial concentration within markets over time. To begin to assess whether those changes have effects on the earnings distribution, I produce scatter plots of changes in mean earnings and changes in industrial concentration. Figure 15 plots several highly aggregated, long-run versions of this relationship. In panel (a), the y-axis shows the change in the log of average earnings across industries within commuting zones between 1976 and 2015, while the x-axis plots the change in the log of the average HHI across industries within commuting zones. Earnings are approximated by dividing total payroll within industry by total employment, both obtained from the LBD. Points are presented in further aggregation as the averages within 20 equal-sized bins.

Over this horizon and at this level of aggregation, there is a clear negative relationship between changes in industrial concentration and changes in earnings. When the same relationship is plotted at the market level (i.e. without first averaging earnings and concentration levels across industries within commuting zones), as in panel (b), the negative relationship remains clear, though the magnitude of the slope of the line of best fit falls by more than 80 percent.

The relationship between industrial concentration and earnings is also sensitive to the time frame considered. Panel (c) plots the same relationship using changes between 2005 and 2015. The relationship remains negative, but the magnitude again declines by more than 75 percent relative to panel (b).

During this time period, earnings can also be calculated using W-2 data. The W-2 earnings measure is conceptually superior to LBD measure, which divides total annual payroll by a point-in-time measure of employment.[10] To the extent that the point-in-time employment measure understates total employment over the course of the year, the LBD average earnings measure overstates true average earnings. Because W-2s are issued to all employees, they capture total annual compensation and total annual employment, allowing me to calculate actual average earnings. In panel (d), I plot this relationship using the W-2 earnings measure. The relationship between changes in earnings and changes in concentration becomes slightly positive, and its magnitude falls again.

These figures all present relationships that are not conditioned on any market character-

[10] The LBD captures employment as of March each year.

istics. Similar relationships also hold in OLS regressions of the form

$$log\left(y_{cit}\right) = log\left(HHI_{cit}\right)\alpha + \delta(c, i, t) + \epsilon_{cit}$$

where, c indexes commuting zones, i indexes industries, t indexes time, $\delta(c, i, t)$ represents a possibly interacted specification of commuting zone, industry, and time fixed effects, and ϵ_{cit} is noise. Estimates from these regressions are reported in Table 1. As in Figure 15, this relationship becomes weaker and ultimately turns slightly positive as I move to my preferred earnings measure in column 3, remaining positive when weights are not used in column 4.[11]

Of course, even conditional on fixed effects or other available observable characteristics of markets, changes in industrial concentration do not necessarily arise exogenously. Indeed, they often arise from other economic changes that also affect the earnings distribution. For example, if a technological breakthrough leads to the emergence of a superstar firm, local concentration could increase as that firm comes to dominate its market. The firm's high productivity could also increase mean earnings in its market. In this scenario, both concentration and earnings increase. A naive assessment could suggest that the increase in concentration caused the increase in earnings, but both were actually caused by a third, unobserved change (the emergence of the high-productivity superstar firm), and the naive estimate is biased.

In order to address concerns like the one just described and estimate the effect of concentration on earnings outcomes, I employ an instrumental variables strategy similar to the one used by Azar et al. (2017). Specifically, I instrument for the HHI in each market (where a market is a commuting zone-level four-digit industry) in each year using the employment-weighted average HHI within the same industry across other commuting zones in the same year. Conceptually, this strategy identifies the effects of local concentration on earnings out-

[11] Qiu and Sojourner (2019) also find a positive relationship between concentration and compensation in OLS regressions.

comes using only variation in local concentration that is driven by broader, non-local forces, as reflected in the "leave one out" concentration mean. Formally, this mean can be written

$$\overline{HHI}_{it}^{-c} = \frac{\sum_{z \neq c} HHI_{zit} \cdot Emp_{zit}}{\sum_{z \neq c} Emp_{zit}}$$

where, c is a specific commuting zone, z indexes commuting zones, i indexes industries, t indexes time, and Emp_{zit} is employment. The first stage regression is

$$log\left(HHI_{cit}\right) = log\left(\overline{HHI}_{it}^{-c}\right)\gamma + \delta(c,i,t) + \eta_{cit}$$

where c now indexes commuting zones, $\delta(c,i,t)$ represents a possibly interacted specification of commuting zone, industry, and time fixed effects, and η_{cit} is noise.

The effects of concentration on earnings outcomes are estimated via

$$log\left(y_{cit}\right) = log\,\widehat{(HHI_{cit})}\beta + \delta(c,i,t) + \varepsilon_{cit}$$

where y_{cit} is an earnings outcome, $\widehat{HHI}_{cit}$ represents fitted values from the first stage regression, and ε_{cit} is noise. Standard errors are clustered at the market level. The coefficient of interest, β, is the elasticity of earnings outcomes y with respect to local industrial concentration. This estimate will reflect the causal effects of local industrial concentration on earnings outcomes if $\overline{HHI}_{it}^{-c}$ predicts HHI_{cit} and only influences earnings outcomes through that channel. As with the trends discussed above, all regressions are weighted by employment unless otherwise noted.[12]

While the exclusion restriction cannot be tested, the relevance of the instrument can be. Table 2 reports estimates from the first stage regression for various configurations of

[12] I do not include time-varying, market-level controls for things like employment levels in my regressions because they are endogenous to the degree of concentration in a market. If, however, an employment control were included, for example, the estimates presented here would be little changed.

commuting zone, industry, and time fixed effects using LBD data from 1976 through 2015. The first column includes no fixed effects and presents the estimate from the univariate regression of the HHI on the instrument. As one might expect based on the construction of the instrument, the coefficient is close to one, indicating a strong positive relationship with local concentration. This relationship survives the introduction of the simplest, non-interacted set of commuting zone, industry, and time fixed effects in the second column.

The third column combines the commuting zone and industry fixed effects into a single "market" fixed effect, and the relationship remains strong. The fourth column increases the flexibility of the time fixed effects by interacting them with the commuting zone fixed effects, to allow for the possibility of trends that differ across regions but have common effects across industries. The coefficient changes little from the third column. Finally, the fifth column adds market-specific linear time trends. The magnitude of the coefficient on the instrument falls by more than 40 percent, but it remains positive and highly statistically significant. Across all columns, the F-statistic associated with the instrument is lowest in the fifth column, and it is still nearly 800.

Table 3 presents the same estimates as Table 2 based only on data from 2005 through 2015. Table 4 also produces these estimates for 2005 through 2015, but limits the sample to markets in which earnings measures based on W-2 data are available. Across the more saturated specifications in columns three through five, the point estimates are smaller in magnitude but exhibit the same pattern as those in Table 2 - whether year fixed effects are interacted with commuting zone fixed effects makes little difference, while adding market trends meaningfully shrinks the first-stage coefficient. These specifications continue to have strong F-statistics in both tables. The fact that the estimates in the second column have turned negative highlights the importance of focusing on within-market variation.

Columns three through five of Tables 2, 3, and 4 present reasonable potential specifications for analyzing the effects of industrial concentration within local labor markets as I have defined them here. My preferred specification, presented in the fourth column of these tables, includes market and commuting zone by year fixed effects. Though the interaction of the commuting zone and year fixed effects makes very little difference in the first stage regressions, that flexibility could be important to some of the reduced form relationships considered below. Although including controls for trends that may vary across markets may be conceptually appealing, the bulk of my analysis relies on W-2 and therefore focuses on 2005 through 2015, and it can be difficult to identify the correct functional form for a trend over a relatively short time period like that. As a result, I prefer not to make the trends specification my default approach, but I do present some results based on it alongside my preferred estimates. In practice, the signs of my estimates are robust to the inclusion of trends, and the magnitudes of the instrumental variables estimates with trends are scaled up relative to the baseline estimates due to the smaller first stage coefficient.

5 Effects of Local Industrial Concentration

I use the instrumental variables strategy described in the previous section to estimate the effects of industrial concentration on a variety of earnings outcomes. I begin with mean earnings, which have been considered in the monopsony context by other recent studies. I also take advantage of the W-2 data to investigate distributional questions, which have gone largely unaddressed thanks in part to the limited availability of data that can measures these outcomes well. Where local labor market circumstances give employers wage-setting power, that power is unlikely to be exercised uniformly over all workers. To the extent industrial concentration corresponds to employer wage-setting power, there is therefore reason to suspect its effects might be experienced differently across the earnings distribution or across groups of workers. In addition to mean earnings, I consider effects of industrial concentration on earnings inequality, both in aggregate and within demographic groups defined by age, gender, race, and educational attainment. Finally, I consider effects of concentration on short- to medium-term earnings mobility.

5.1 Earnings and Inequality

Table 5 reports estimates of the effects of industrial concentration on average earnings using various versions of my preferred specification.[13] The dependent variable is the log of mean earnings, either constructed from the total payroll and employment employment variables in the LBD or calculated from Form W-2 data, as indicated. As mentioned above, the reported coefficients are elasticities of earnings with respect to local industrial concentration. In the first column, which uses LBD data from 1976 through 2015, the elasticity is about -0.05 and statistically significant. To put this estimate in context, Figure 8 indicates that moving between the local HHI experienced by the median workers and either the 25th or 75th percentile in 2015 represented approximately a threefold change in industrial concentration. This elasticity implies that the move from the median up to the 75th percentile would reduce earnings by about 15 percent, while the move down to the 25th percentile would increase them by a similar amount.

Column 2 repeats this analysis using only data from 2005 through 2015. The earnings effect declines in magnitude to just under -0.01 and loses statistical significance when estimated within this shorter period. Switching to the conceptually superior W-2 earnings measure in Column 3 increases its magnitude again to just over -0.03, and it returns to statistical significance. This elasticity implies that the move from the local HHI experienced by the median worker up to the 75th percentile would reduce earnings by nearly ten percent. These estimates are broadly consistent with other recent findings on the effects of labor market concentration on earnings (e.g. Azar et al. 2017; Benmelech et al. 2018).

[13] Tables reporting estimates of the effect of concentration on earnings and inequality outcomes using all variations on the specifications reported in Tables 2 through 4 are available in the Online Appendix.

Column 4 again repeats the analysis of the W-2 earnings measure without weighting markets according to employment. The unweighted estimate is more than three times larger in magnitude than the weighted estimate. This suggests that the effects of concentration on earnings may be larger in smaller markets, as the overall average effect becomes larger when smaller markets are given greater relative weight.

Next, I consider the effects of local industrial concentration on earnings inequality. Dependent variables are constructed within local labor markets from W-2 data. In Table 6, I report estimates of the effects on key earnings percentile ratios (90/10, 50/10, and 90/50), as well as the Gini coefficient using my preferred specification. First, in column 1, higher local industrial concentration increases the 90/10 earnings ratio; the elasticity is 0.17. I next estimate effects on the 50/10 and 90/50 earnings ratios (columns 2 and 3, respectively) to get a sense of whether the overall inequality effect is driven by changes in the top or the bottom of the distribution. The relative magnitude of the coefficients from these regressions indicates that the changes in the bottom of the distribution account for about 60 percent of the increase in the 90/10 ratio; the elasticity of the 50/10 ratio is about 0.11, while the elasticity of the 90/50 ratio is just under 0.07.

Changes in earnings percentile ratios indicate that increases in concentration reduce earnings at the bottom of the distribution relative to earnings in the middle and at the top. They do not, however, reveal how earnings change in absolute terms across the distribution. The first estimates in this section show that mean earnings fall, so some portion of the distribution must see negative effects, but it is also conceivable that some regions of the distribution could see earnings increase. If monopsony rents accrue to some employees in form of, for example, bonuses to top managers, values of high percentiles of the earnings distribution could increase with concentration.

Figure 16 presents the effects of local industrial concentration on key percentiles of the earnings distribution, estimated using my preferred specification.[14] These estimates show that the increases in inequality revealed by the percentile ratios are driven entirely by declining values of low percentiles, not increasing values of high percentiles. Changes in the 75th and 90th percentiles are not distinguishable from zero, so the increases in the 90/50 and 90/10 ratios arise almost entirely from reductions in the values of the median and the 10th percentile of the earnings distribution. Both these estimates and the percentile ratio estimates above are consistent with Webber (2015)'s individual-level unconditional quantile regression estimates.

I also consider the effect of concentration on the Gini coefficient, another commonly used measure of inequality, in column 4 of Table 6. I again find that increased concentration leads to increased inequality.

One caveat to this analysis is that the exclusion restriction discussed above may be susceptible to violation by local shocks that affect concentration and earnings outcomes across an entire industry. One might expect such shocks to be less common in the non-tradable sector, where production and provision of goods and services are more directly tied to local conditions. When I reproduce my main estimates using only industries classified as non-tradable or construction by Mian and Sufi (2014), results are similar to the baseline estimates discussed here, though somewhat larger in magnitude. These estimates are reported in Table

5.2 Effects by Demographic Group

In addition to varying across the earnings distribution, the effects of concentration may also vary across groups of workers defined by demographic characteristics. Summary measures of labor market conditions like the unemployment rate differ systematically across groups

[14]Tabular versions of these estimates, as well as other estimates reported in figure form, can be found in Appendix C. That section also contains reduced for estimates for all specifications discussed here.

defined by age, race, sex, and education, both in levels and in changes over the business cycle. To the extent that such measures reflect systematic, pre-existing differences in employment opportunities across groups, changes in local industrial concentration may "treat" those groups with different intensities and have different effects on their earnings outcomes.

Figure 17 plots the effects of local industrial concentration on mean earnings by demographic groups based on my preferred specification. Estimates indicate that men, younger workers, and white workers experience more negative earnings effects than do women, prime-age and older workers, or Black workers. The earnings effect for women is in fact positive. High and low education workers experience similar earnings effects, though those estimates come with the caveat that they are based on far fewer individual observations, as education information is available only for individuals who responded to the ACS between 2005 and 2015.

Turning to inequality outcomes, all groups of workers see statistically significant increases in the 90/10 earnings ratio due to increased local industrial concentration, as shown in Figure 18. Point estimates are larger for men, older workers, and those with a high school diploma or less. Considering changes in the 50/10 earnings ratio (panel (a) of Figure 19) alongside changes in the 90/50 earnings ratio (panel (b)) shows that, like in the full sample, the inequality increases experienced by men, older workers, white workers, Hispanic workers, and members of both education groups are driven mostly by changes in the bottom half of the earnings distribution. Women, young workers, and Black workers, on the other hand, see changes in the top of the earnings distribution account for most of the increase in inequality they experience.

When inequality is instead measured using the Gini coefficient (Figure 20) fewer groups experience an increase, and some patterns within demographic categories differ. For instance, there is a substantial difference between the Gini elasticities of high and low education workers, with low education workers experiencing increased inequality as a result of increased concentration, while the point estimate for high education workers is negative and not statistically significant. Also, the age gradient in the inequality effect is reversed when measured using the Gini coefficient instead of the 90/10 earnings ratio. Younger workers have the largest Gini elasticity, while the 90/10 earnings ratio was most responsive to concentration changes for older workers.

5.3 Earnings Mobility

Finally, I consider the effects of local industrial concentration on short- to medium-term earnings mobility. Job-switching is an important channel for getting a raise (Topel and Ward, 1992; Fallick et al., 2012; Molloy et al., 2014), and if less competition among employers leads to fewer opportunities for workers to switch jobs, increases in concentration could limit their ability to move up the earnings distribution.

I consider measures of both relative and absolute earnings mobility over horizons extending up to five years. For each market m, my measure of relative earnings mobility for year t over the following N years is the coefficient from a regression of each worker in market m in year t's percentile rank in the national earnings distribution in year $t+N$ on their percentile rank in the national earnings distribution in year t. I refer to this as the rank-rank coefficient. A higher rank-rank coefficient means that one's present position in the earnings distribution is more predictive of one's future position, so a market with a higher rank-rank coefficient has lower relative earnings mobility. My measure of absolute earnings mobility is the mean difference in log earnings between years t and $t+N$ for workers in market m in year t. In constructing both of these measures, I use only workers who are between 25 and 54 in both year t and year $t+N$ in order to minimize instances of workers leaving the labor force without conditioning on continued attachment. I do not require that workers continue to live in their base-year commuting zone or work in their base-year industry. Because both the rank-rank coefficient and the mean difference in log earnings can be negative, they enter the regressions in levels and the coefficients the regressions produce represent semi-elasticities.

Figure 21 plots the effects of local industrial concentration on the rank-rank coefficient over horizons extending up to five years, using my preferred specification in panel (a). Looking one year ahead, the effect is small and negative, but over two years it is positive and it grows in magnitude up to four years out before shrinking slightly over the five-year horizon. The generally positive effect on the rank-rank coefficient suggests that over most of these horizons increased concentration reduces relative earnings mobility. However, these results are less robust to reasonable alternative specifications than the main earnings and inequality results discussed above are. Panel (b) shows estimates from the specification that includes market-specific linear trends. Here, the point estimates are shifted down relative to the preferred specification, and the 95 percent confidence intervals are generally at least twice as wide. The one-year effect is more clearly negative, and only the five-year effect is statistically significantly positive.

Effects on absolute earnings mobility are similarly sensitive to specification. Estimates from my preferred specification show increased concentration leading to faster earnings growth over longer horizons in panel (a) of Figure 22, but once market trends are added to these regressions in panel (b), the estimates for longer horizons in particular decline substantially in magnitude and become statistically indistinguishable from zero.

Overall, the fact that earnings mobility effects are sensitive to specification changes while other earnings and inequality effects are generally robust to them cautions against drawing strong conclusions about the effects of local industrial concentration on earnings mobility here.

6 Discussion and Conclusion

This paper's finding that increased local labor market concentration reduces earnings is consistent with other recent findings from online job boards (Azar et al., 2017) and the manufacturing sector (Benmelech et al., 2018). My estimates of the effects of concentration on inequality are consistent with Webber (2015): when concentration increases, the gap between the top of the distribution and the middle of the distribution widens not because earnings increase at the top but because they decline in the middle. The gap between the middle and the bottom increases by more because earnings fall more at the bottom than they do in the middle. To the extent that employers in more concentrated markets have more power over workers, these estimates provide some evidence that that power may contribute to increased inequality, as the Council of Economic Advisers (2016b) suggested it might.

However, these estimates, combined with the fact that local industrial concentration has declined since 1976 indicate that it has not contributed to the increase in inequality over that period. Back-of-the-envelope calculations suggest that the average within-market 90/10 earnings ratio was 6.3 percent lower and average annual earnings were 1.2 percent higher in 2015 than they would have been if average local industrial concentration had been at its 1976 level, which was about 36 percent higher. For context, the national 90/10 ratio increased by about 40 percent between 1976 and 2015, while average annual earnings increased by about 30 percent in real terms for prime-age workers over that period.[15] Changes in concentration appear to have modestly mitigated the trend toward increased inequality rather than contributing to it.

The subgroup analyses in Section 5.2 suggest that the effects of local labor market con-

[15] The change in the 90/10 ratio is calculated from estimates in Proctor et al. (2016). The change in average annual earnings is estimated using publicly available microdata from the Annual Social and Economic Supplement to the Current Population Survey. The sample includes workers between ages 25 and 54 with positive earnings in the 1977 and 2016 surveys. Estimates are adjusted for inflation using the CPI-U-RS. The 1977 topcode is applied, in real terms, to the 2016 data before earnings are estimated.

centration may vary not only across the distribution of earnings but also across demographic groups. While all groups experience increases in inequality as measured using the 90/10 earnings ratio due to increase concentration, not all groups see mean earnings decline. In particular, women see an earnings increase, and the point estimate of the earnings effect for black workers is positive, though small and not statistically significant. Notably, both groups have historically experienced labor market discrimination in the United States. Previous research has considered the interaction between monopsony power and so-called taste-based discrimination (e.g. Hirsch et al., 2010; Hirsch and Jahn, 2015; Webber, 2016; Fanfani, 2018), and changes in related employment dynamics could rationalize positive earnings effects for these groups.

Two plausible explanations arise from possible changes in the composition of employees and employers, respectively. If industrial concentration is a reasonable proxy for employers' monopsony power in the labor market, then increasing concentration could allow firms to be more selective in their hiring processes. Firms may choose to exercise that power by not hiring marginal workers from some demographic groups rather than hiring them and suppressing their wages. If inframarginal workers in those groups are higher earners, average earnings could increase mechanically as concentration increases and lower-earning marginal workers are excluded. The composition of workers could also change if demographic groups are differentially exposed to changes in skill requirements associated with increased labor market concentration (Hershbein and Macaluso, 2018). Alternatively, if the composition of employers shifts toward larger firms with more established human resources practices as concentration increases, workers in these groups could benefit from institutional safeguards against pay discrimination, large-firm wage premiums, or other differences in business practices between incumbent and entrant firms. There is some evidence of a wage premium associated with modern chain retailers (Cardiff-Hicks et al., 2015). If the entry of such firms contributes to increased concentration, the associated wage premium could lead to positive effects on concentration on earnings, including in groups that commonly experience discrimination.

Beyond the context of discrimination, there are open and interesting questions about the role of changes in the distribution of firm size in realizing the effects of labor market concentration. Changes in how workers sort across firms are also potentially relevant here. These topics should receive additional attention in future work.

The effects of concentration on percentile earnings ratios for Black workers and women also differ from the aggregate pattern. For all workers, about 60 percent of the increase in the 90/10 ratio due to increases in concentration is realized below the median, but for Black workers and women, essentially the entire increase is realized above the median. Some of this could be attributable to the fact that any given percentile of the Black or female earnings distribution generally has a lower value than that same percentile in the overall distribution. Changes that affect any given point in the overall distribution therefore reach further up the distributions within these groups. Of course, other, non-mechanical factors could also be playing a role here, and further research on differential consequences of increased labor market concentration across groups of workers would be valuable.

While industrial concentration is not a perfect measure of labor market concentration, the consistency between these estimates and others based on occupation suggest that it is a useful tool for thinking about this concept. Prior to this work, little was known about how industrial concentration had changed over time at the local level. The divergence between national and local trends in industrial concentration discussed in Section 3 highlights the importance of thinking about concentration at the local level.[16] While both the trends and the regression estimates presented above are generally robust to alternative definitions of

[16] In contemporaneous work, Rossi-Hansberg et al. (2018) find that national and local product market concentration trends also diverge. Using different geographic and industrial levels of aggregation than those employed here, Lipsius (2018) also finds diverging trends in local and national labor market concentration that are similar to those presented here.

local labor markets, additional work on understanding the reasonable sets of alternative employment opportunities for workers and potential employees for firms could help improve our understanding of what constitutes a local labor market and how changes in conditions within certain industries, occupations, or localities might have consequences in others.

The importance of thinking about labor market concentration locally extends to lightly populated localities. The employment-weighted local industrial concentration distribution has a long right tail, even as it has been tightening for decades. On top of that, evidence from unweighted estimates suggests that the effects of concentration on earnings outcomes may be more negative in smaller markets. Future research should specifically dig more deeply into these markets where the consequences of increased concentration may be experienced more intensely by a smaller number of people with fewer alternative employment opportunities.

Of course, industrial concentration is not identically equal to labor market concentration, and even if that were a concept that could be measured perfectly, it would only serve as a proxy for monopsony power. Any given strategy will have limitations. Researchers should continue to compare estimates based on alternative approaches to identifying employer wage-setting power.

The data used in this paper also have their limitations, even as they represent some of the best available tools for considering these questions. For example, the earnings measures I focus on here are constructed from Form W-2. This form reports only wage and salary earnings. Moreover, it reports earnings from only a specific type of work arrangement. Income earned through independent contracting or as profit from a business is not captured by these data. The inability to measure business income could make it difficult to identify the amount and recipients of monopsony rents. Researchers should work to incorporate measures of additional sources of income into future work, including sources relevant to both individuals who exert monopsony power and those seeking alternatives to employment in markets that are dominated by it.

References

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2017). The fall of the labor share and the rise of superstar firms. NBER Working Paper No. 23396.

Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2018). Concentration in US labor markets: Evidence from online vacancy data. NBER Working Paper No. 24395.

Azar, J., Marinescu, I., and Steinbaum, M. I. (2017). Labor market concentration. NBER Working Paper No. 24147.

Benmelech, E., Bergman, N., and Kim, H. (2018). Strong employers and weak employees: How does employer concentration affect wages? NBER Working Paper No. 24307.

Boal, W. M. (1995). Testing for employer monopsony in turn-of-the-century coal mining. *The RAND Journal of Economics*, 26(3):519–536.

Cardiff-Hicks, B., Lafontaine, F., and Shaw, K. (2015). Do large modern retailers pay premium wages? *ILR Review*, 68(3):633–665.

Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R., and Narang, J. (2016). The fading American Dream: Trends in absolute income mobility since 1940. NBER Working Paper No. 22910.

Council of Economic Advisers (2016a). Benefits of competition and indicators of market power. https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160414_cea_competition_issue_brief.pdf.

Council of Economic Advisers (2016b). Labor market monopsony: Trends, consequences, and policy responses. https://obamawhitehouse.archives.gov/sites/default/files/page/files/20161025_monopsony_labor_mrkt_cea.pdf.

De Loecker, J. and Eeckhout, J. (2017). The rise of market power and the macroeconomic implications. NBER Working Paper No. 23687.

Dube, A., Giuliano, L., and Leonard, J. S. (2018a). Fairness and frictions: The impact of unequal raises on quit behavior. NBER Working Paper No. 24906.

Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2018b). Monopsony in online labor markets. NBER Working Paper No. 24416.

Dube, A., Manning, A., and Naidu, S. (2017). Monopsony and employer mis-optimization account for round number bunching in the wage distribution. Working paper.

Edmond, C., Midrigan, V., and Xu, D. Y. (2018). How costly are markups? NBER Working Paper No. 24800.

Falch, T. (2010). The elasticity of labor supply at the establishment level. *Journal of Labor Economics*, 28(2):237–266.

Fallick, B., Haltiwanger, J., and McEntarfer, E. (2012). Job-to-job flows and the consequences of job separations. FEDS Working Paper No. 2012-73.

Fanfani, B. (2018). Tastes for discrimination in monopsonistic labour markets. Working Paper.

Fort, T. C. and Klimek, S. D. (2018). The effects of industry classification changes on us employment composition. CES Working Paper 18-28.

Fox, J. T. (2010). Estimating the employer-switching costs and wage responses of forward-looking engineers. *Journal of Labor Economics*, 28(2):357–412.

Grullon, G., Larkin, Y., and Michaely, R. (2018). Are U.S. industries becoming more concentrated? Working paper, available https://dx.doi.org/10.2139/ssrn.2612047.

Hall, R. E. (2018). New evidence on the markup of prices over marginal costs and the role of mega-firms in the US economy. NBER Working Paper No. 24574.

Hershbein, B. and Macaluso, C. (2018). Labor market concentration and the demand for skills. Working paper.

Hirsch, B. and Jahn, E. J. (2015). Is there monopsonistic discrimination against immigrants? *ILR Review*, 68(3):501–528.

Hirsch, B., Schank, T., and Schnabel, C. (2010). Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from Germany. *Journal of Labor Economics*, 28(2):291–330.

Humphreys, B. R. and Pyun, H. (2015). Monopsony exploitation in professional sport: Evidence from Major League Baseball position players, 2000–2011. West Virginia University Department of Economics Working Paper 15-48.

Jarmin, R. S. and Miranda, J. (2002). The Logitudinal Business Database. https://www.census.gov/ces/pdf/CES-WP-02-17.pdf.

Krueger, A. B. and Ashenfelter, O. (2018). Theory and evidence on employer collusion in the franchise sector. IZA Discussion Paper No. 11672.

Landon, J. H. and Baird, R. N. (1971). Monopsony in the market for public school teachers. *American Economic Review*, 61(5):966–971.

Lipsius, B. (2018). Labor market concentration does not explain the falling labor share. https://ideas.repec.org/p/jmp/jm2018/pli1202.html.

Luizer, J. and Thornton, R. (1986). Concentration in the labor market for public school teachers. *ILR Review*, 39(4):573–584.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.

Manning, A. (2011). Imperfect competition in the labor market. *Handbook of Labor Economics*, 4(Part B):973–1041.

Marinescu, I. and Rathelot, R. (2018). Mismatch unemployment and the geography of job search. *American Econonomic Journal: Macroeconomics*, 10(3):42–70.

Matsudaira, J. D. (2014). Monopsony in the low-wage labor market? evidence from minimum nurse staffing regulations. *Review of Economics and Statistics*, 96(1):92–102.

Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Molloy, R., Smith, C. L., and Wozniak, A. K. (2014). Declining migration within the U.S.: The role of the labor market. NBER Working Paper No. 20065.

Muehlemann, S., Ryan, P., and Wolter, S. C. (2013). Monopsony power, pay structure, and training. *ILR Review*, 66(5):1097–1114.

Naidu, S. (2010). Recruitment restrictions and labor markets: Evidence from the postbellum U.S. South. *Journal of Labor Economics*, 28(2):413–445.

Pinheiro, J. (1987). Atd̂ivestiture & the telecommunications market. *Berkeley Technology Law Journal*, 2(2):303–355.

Priest, G. L. (2010). Timing "disturbances" in labor market contracting: Roth's findings and the effects of labor market monopsony. *Journal of Labor Economics*, 28(2):447–472.

Proctor, B. D., Semega, J. L., and Kollar, M. A. (2016). Income and poverty in the United States: 2015. Technical report, U.S. Census Bureau.

Qiu, Y. and Sojourner, A. J. (2019). Labor market concentration and labor compensation. http://dx.doi.org/10.2139/ssrn.3312197.

Ransom, M. R. and Oaxaca, R. L. (2010). New market power models and sex differences in pay. *Journal of Labor Economics*, 28(2):267–289.

Ransom, M. R. and Sims, D. P. (2010). Estimating the firm's labor supply curve in a "New Monopsony" framework: Schoolteachers in Missouri. *Journal of Labor Economics*, 28(2):331–355.

Robinson, J. (1933). *The Economics of Imperfect Competition*. Macmillan, London.

Rossi-Hansberg, E., Sarte, P.-D., and Trachter, N. (2018). Diverging trends in national and local concentration. NBER Working Paper No. 25066.

Sinclair, T. M. (2014). Where people search for jobs: Cross-border labor mobility report. Indeed Hiring Lab.

Smith, A. (1776). An inquiry into the nature and causes of the wealth of nations.

Staiger, D. O., Spetz, J., and Phibbs, C. S. (2010). Is there monopsony in the labor market? evidence from a natural experiment. *Journal of Labor Economics*, 28(2):211–236.

Starr, E., Frake, J., and Agarwal, R. (2018). Mobility constraint externalities. Working paper.

Topel, R. H. and Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, 107(2):439–479.

Traina, J. (2018). Is aggregate market power increasing? production trends using financial statements. Working Paper, available https://dx.doi.org/10.2139/ssrn.3120849.

Tucker, L. (2017). Monopsony for whom? evidence from brazilian administrative data. Working paper, available http://leetucker.net/docs/LeeTucker_JMP_latest.pdf.

U.S. Department of the Treasury (2015). Non-compete contracts: Economic effects and policy implications. https://www.treasury.gov/resource-center/economic-policy/Documents/UST%20Non-competes%20Report.pdf.

Wagner, D. and Layne, M. (2014). The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications' (CARRA) Record Linkage Software. CARRA Working Paper No. 2014-01.

Webber, D. A. (2015). Firm market power and the earnings distribution. *Labour Economics*, 35:123–134.

Webber, D. A. (2016). Firm-level monopsony and the gender pay gap. *Industrial Relations: A Journal of Economy and Society*, 55(2):323–345.

Figures

Figure 1: National Industrial Concentration Trend

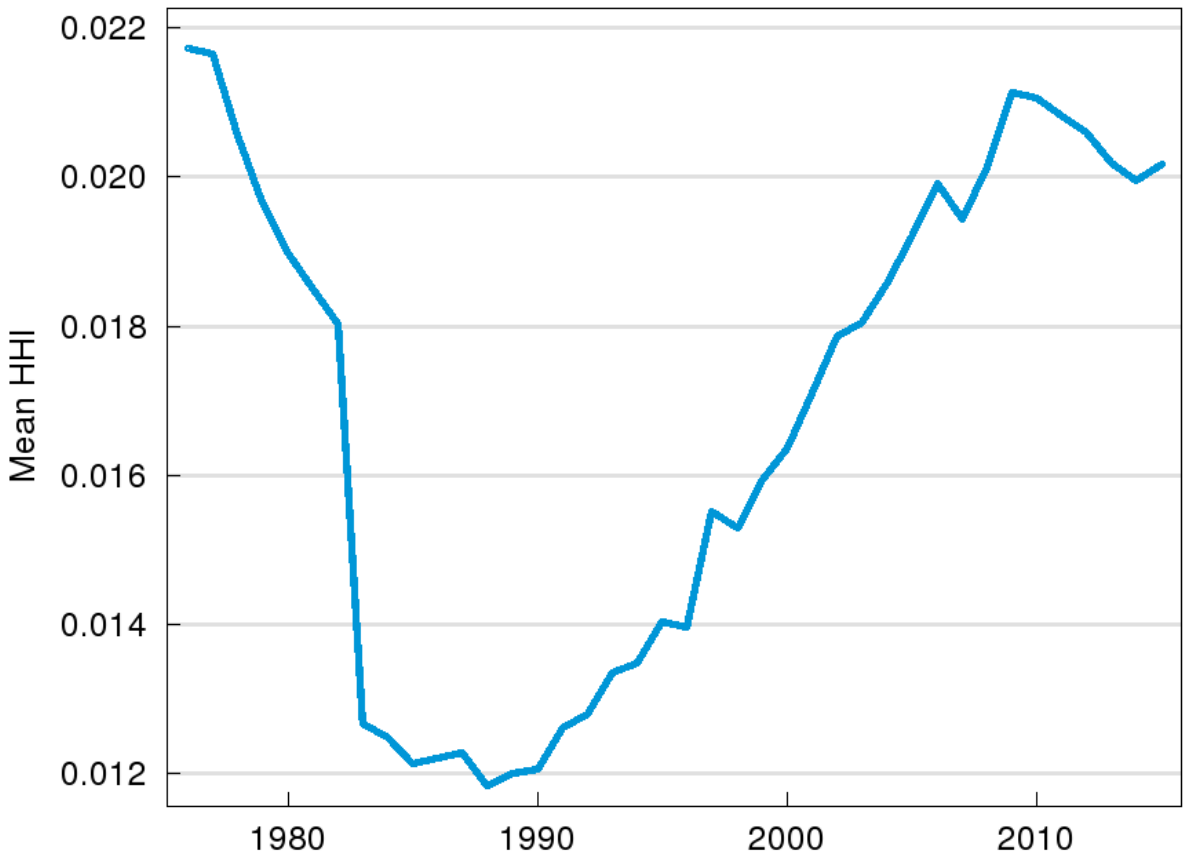

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots the mean Herfindahl-Hirschman Index across national four-digit NAICS industries, standardized according to Fort and Klimek (2018), for each year from 1976 through 2015. Means are calculated using total market employment as weights.

Figure 2: Local Industrial Concentration Trend

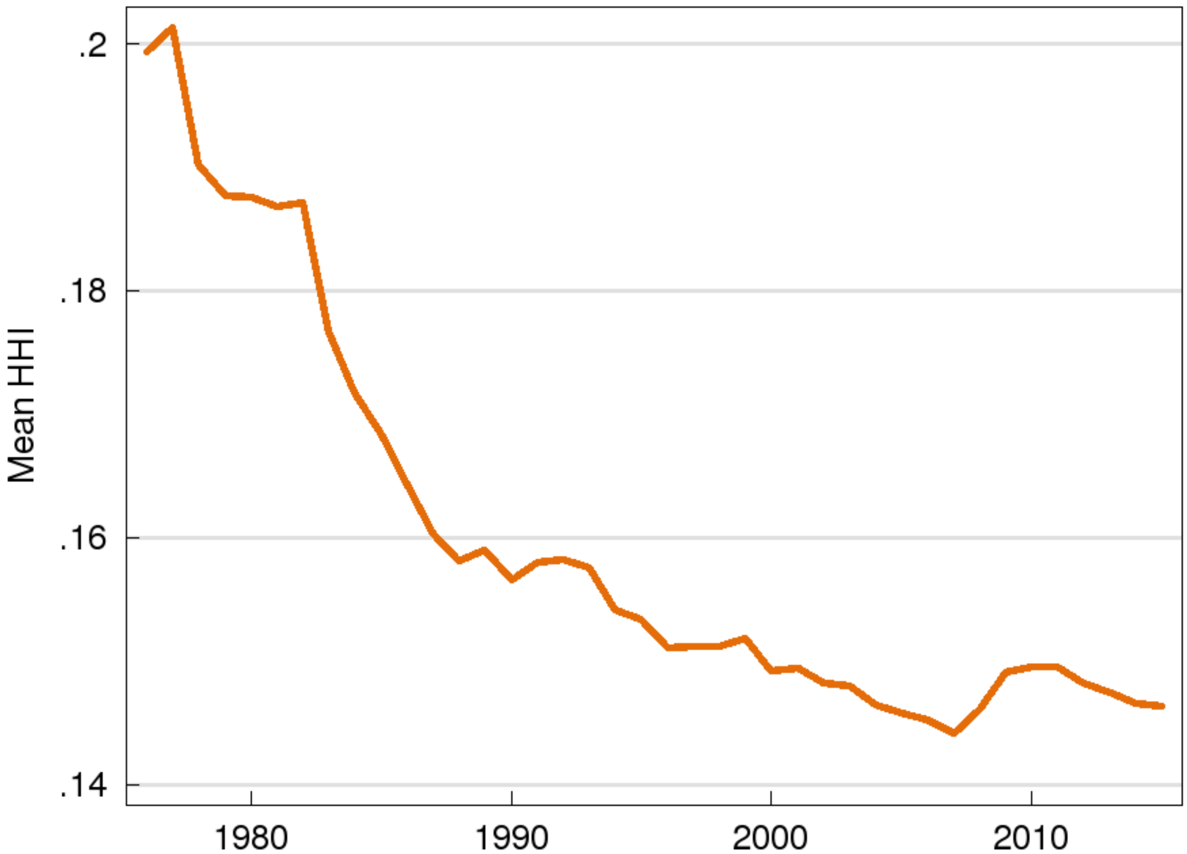

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots the mean Herfindahl-Hirschman Index across commuting zone-level four-digit NAICS industries, standardized according to Fort and Klimek (2018), for each year from 1976 through 2015. Means are calculated using total market employment as weights.

Figure 3: Decomposition of National Industrial Concentration Trend

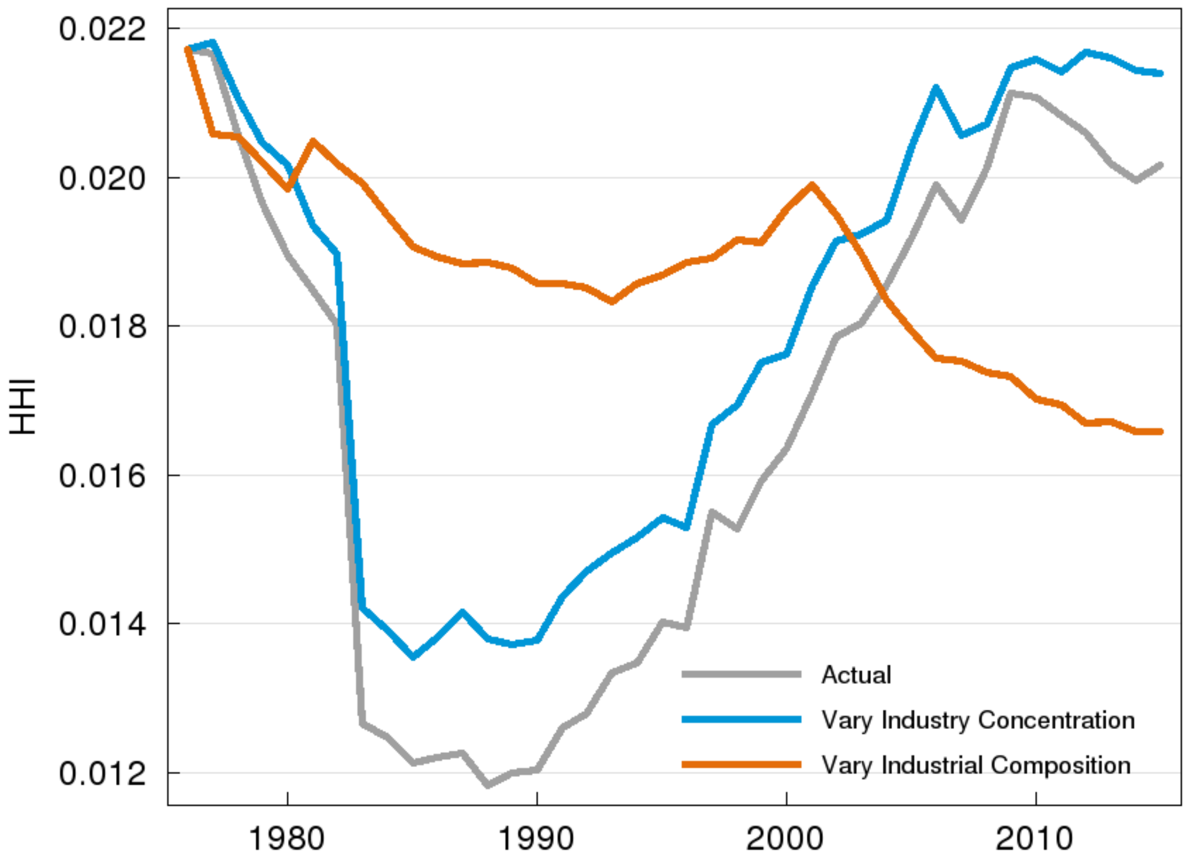

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots the mean Herfindahl-Hirschman Index across national four-digit NAICS industries, standardized according to Fort and Klimek (2018), for each year from 1976 through 2015, as well as counterfactual versions of that trend that would have been observed if different components of the average were allowed to vary in isolation. The gray line plots the actual observed trend in national industrial concentration. The blue line plots the trend that would have been observed if only industrial concentration is allowed to vary (i.e. if industrial composition is held constant at 1976 shares). The orange line plots the trend that would have been observed if only industrial composition is allowed to vary (i.e. industrial concetration is held constant at 1976 levels).

Figure 4: Decomposition of Local Industrial Concentration Trend

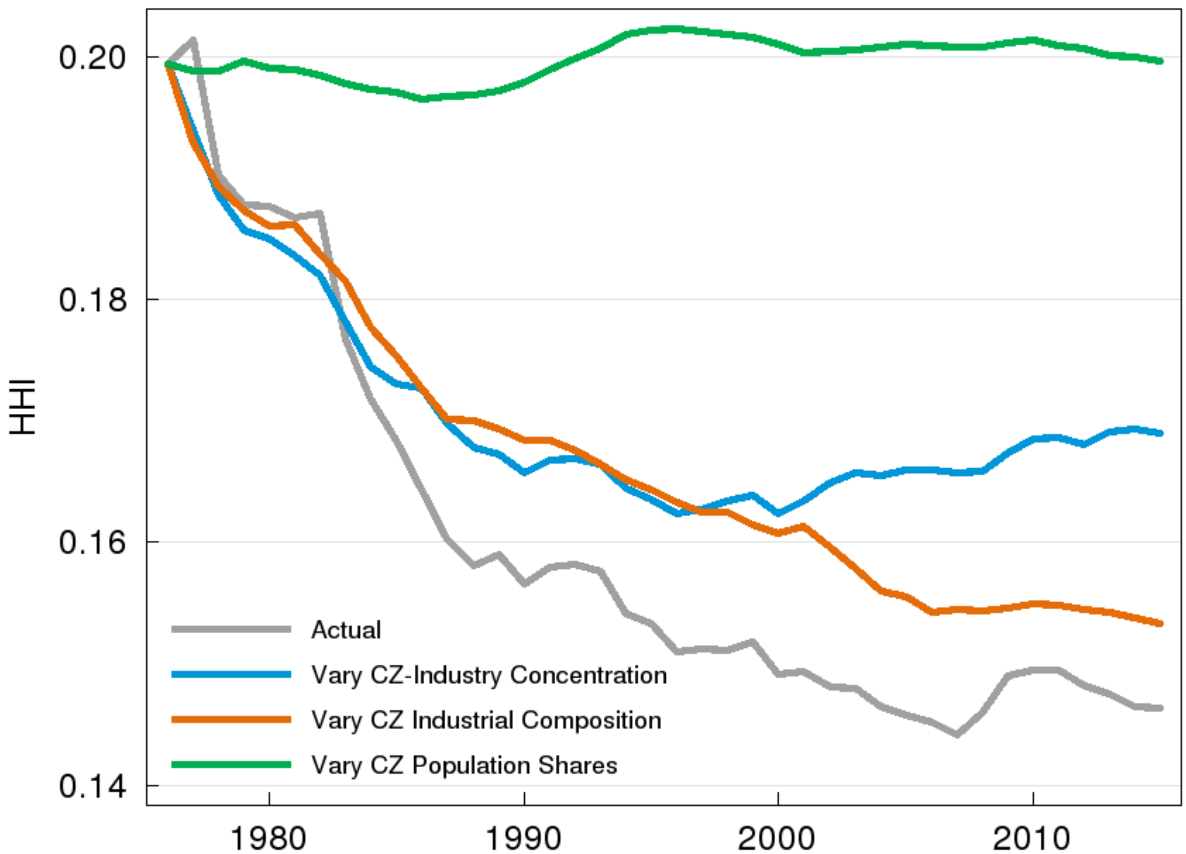

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots the mean Herfindahl-Hirschman Index across commuting zone-level four-digit NAICS industries, standardized according to Fort and Klimek (2018), for each year from 1976 through 2015, as well as counterfactual versions of that trend that would have been observed if different components of the average were allowed to vary in isolation. The gray line plots the actual observed trend in national industrial concentration. The blue line plots the trend that would have been observed if only industrial concentration is allowed to vary (i.e. if industrial composition and the distribution of employment across commuting zones are held constant at 1976 levels). The orange line plots the trend that would have been observed if only industrial composition is allowed to vary (i.e. industrial concetration and the distribution of employment across commuting zonesares held constant at 1976 levels). The green line plots the trend that would have been observed if the distribution of employment across commuting zones is allowed to vary (i.e. industrial composition and industrial concentration are held constant at 1976 levels).

Figure 5: Local Industrial Concentration Trends under Alternative Assumptions

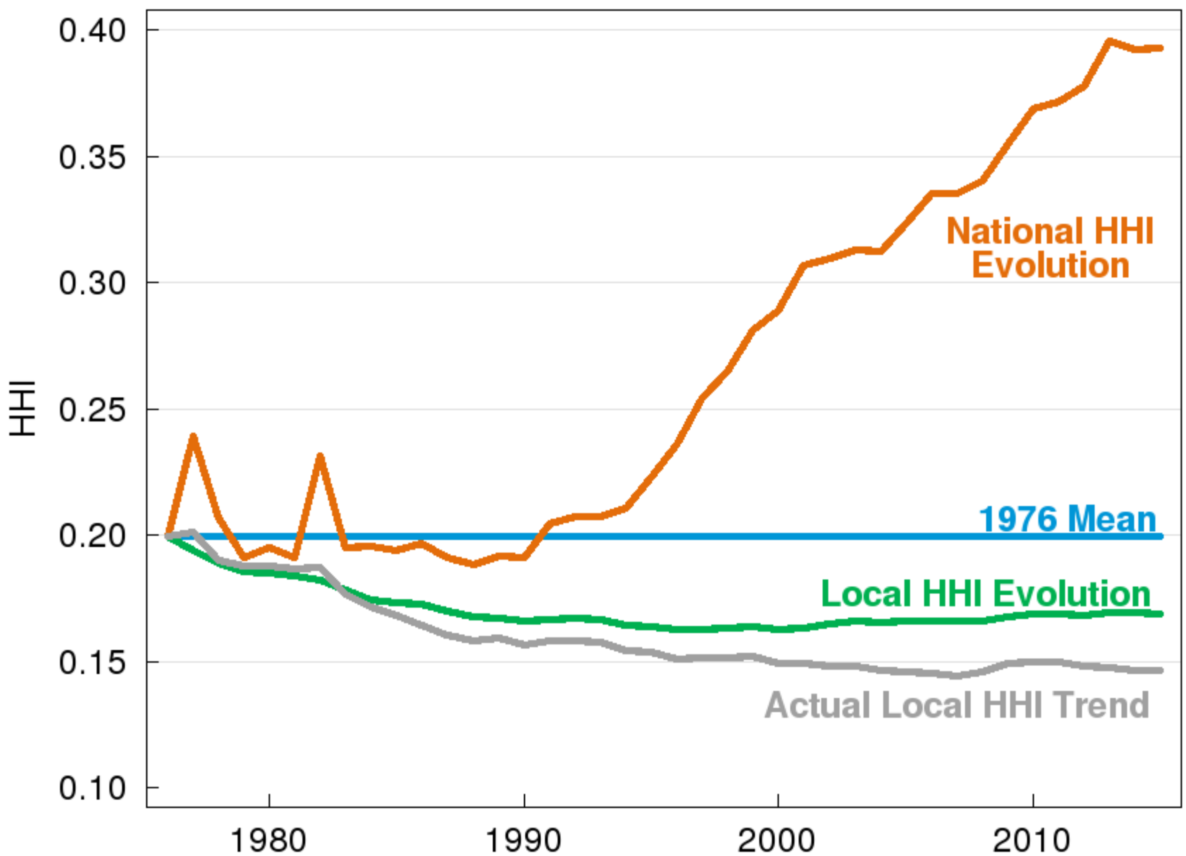

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots trends in local industrial concentration under alternative assumptions about how commuting zone-level four-digit NAICS industry HHIs change over time. The gray line plots the actual trend in average local industrial concentration. The blue line plots the 1976 value of that measure. The orange line plots the local concentration trend that would have been observed if local industry HHIs had evolved proportionally to the national HHIs in the same industries. The green line plots the local concetration trend that would have been observed if local industry HHIs had evolved as they actually did but local industrial composition had remained fixed at 1976 shares.

Figure 6: Markets with At Least One Top-5 Firm

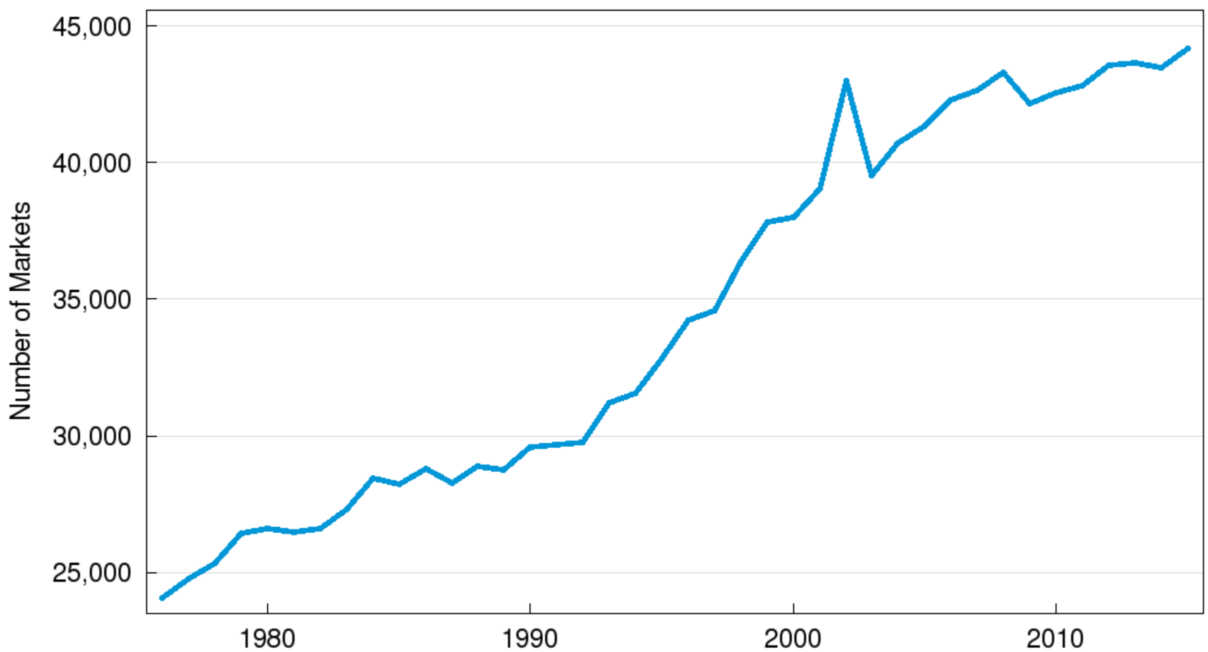

Source: Longitudinal Business Database, 1976–2015
Note: Figure reports the number of markets (commuting zone-level four-digits NAICS industries) that contain at least one establishment belonging to at least one of the five largest firms by national employment within that four-digit NAICS industry.

Figure 7: Share of Markets with Multiple Top-5 Firms

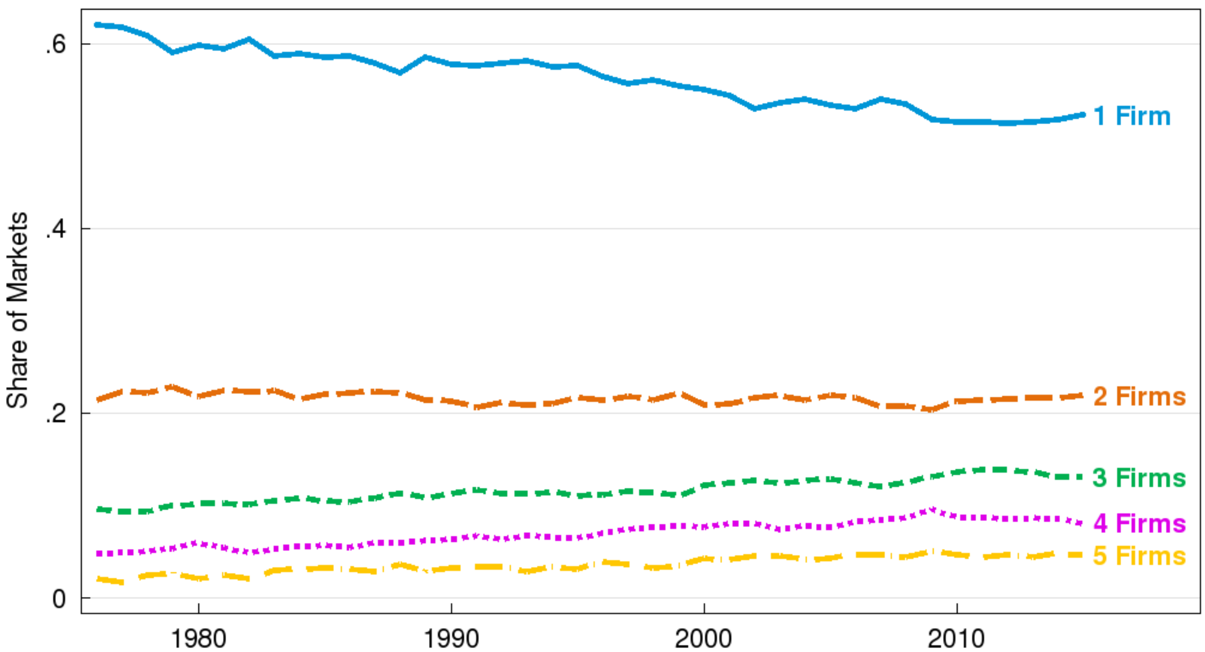

Source: Longitudinal Business Database, 1976–2015
Note: Firm reports the share of markets (commuting zone-level four-digits NAICS industries) containing at least N top-five national firms, conditional on containing at least one such firm.

Figure 8: Distributional Trends in Local Industrial Concentration

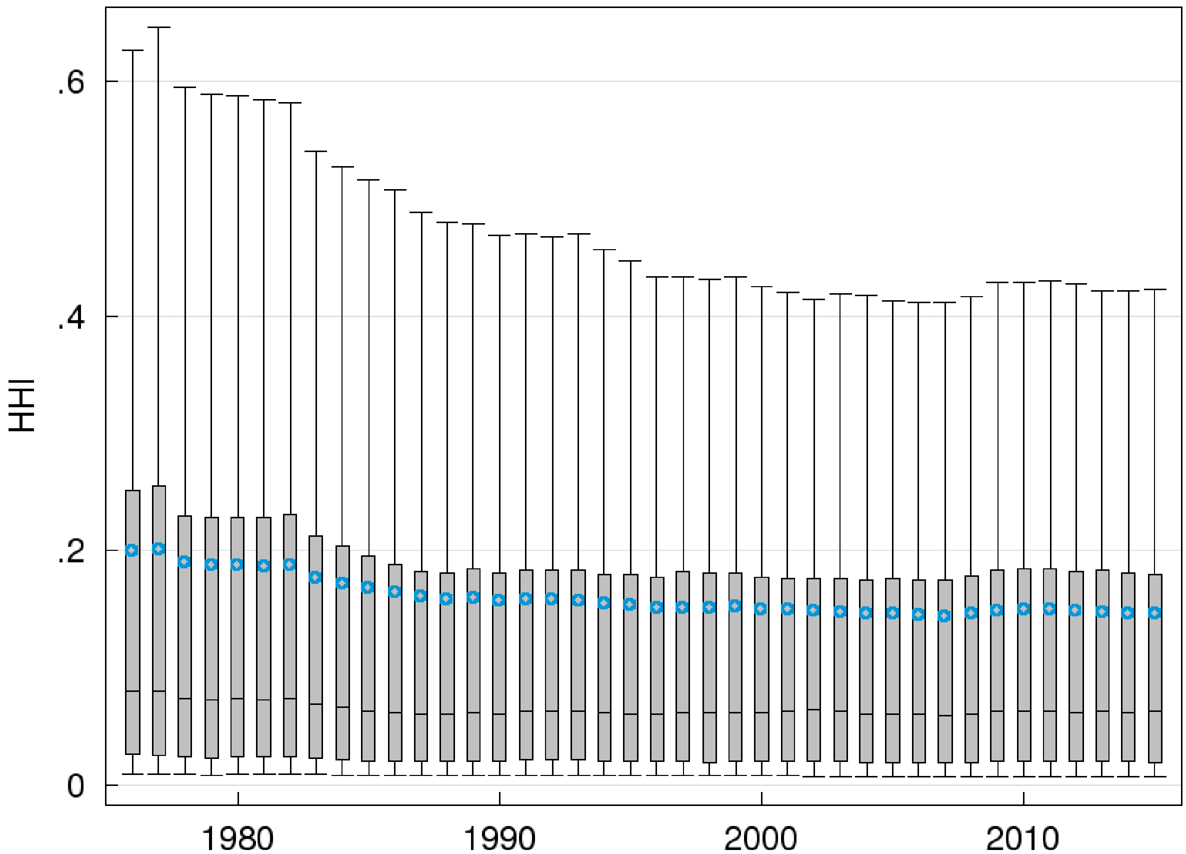

Source: Longitudinal Business Database, 1976–2015
Note: Figure plots trends in the mean and key percentiles of the local industrial concentration distribution, as measured using the Herfindahl-Hirschman Index. The unit of analysis is the commuting zone-level four-digit NAICS industry. The blue circles represent the mean. The boundaries of the box in the box and whisker plots represent the 25th and 75th percentiles of the distribution, while the whiskers represent the 10th and 90th percentiles. Percentiles are approximated using the mean value of markets surrounding the actual percentile value. Percentile values are the mean value for markets within a given percentile. All values are calculated using total market employment as weights.

Figure 9: Average Concentration Across Industries within Commuting Zones, 1976

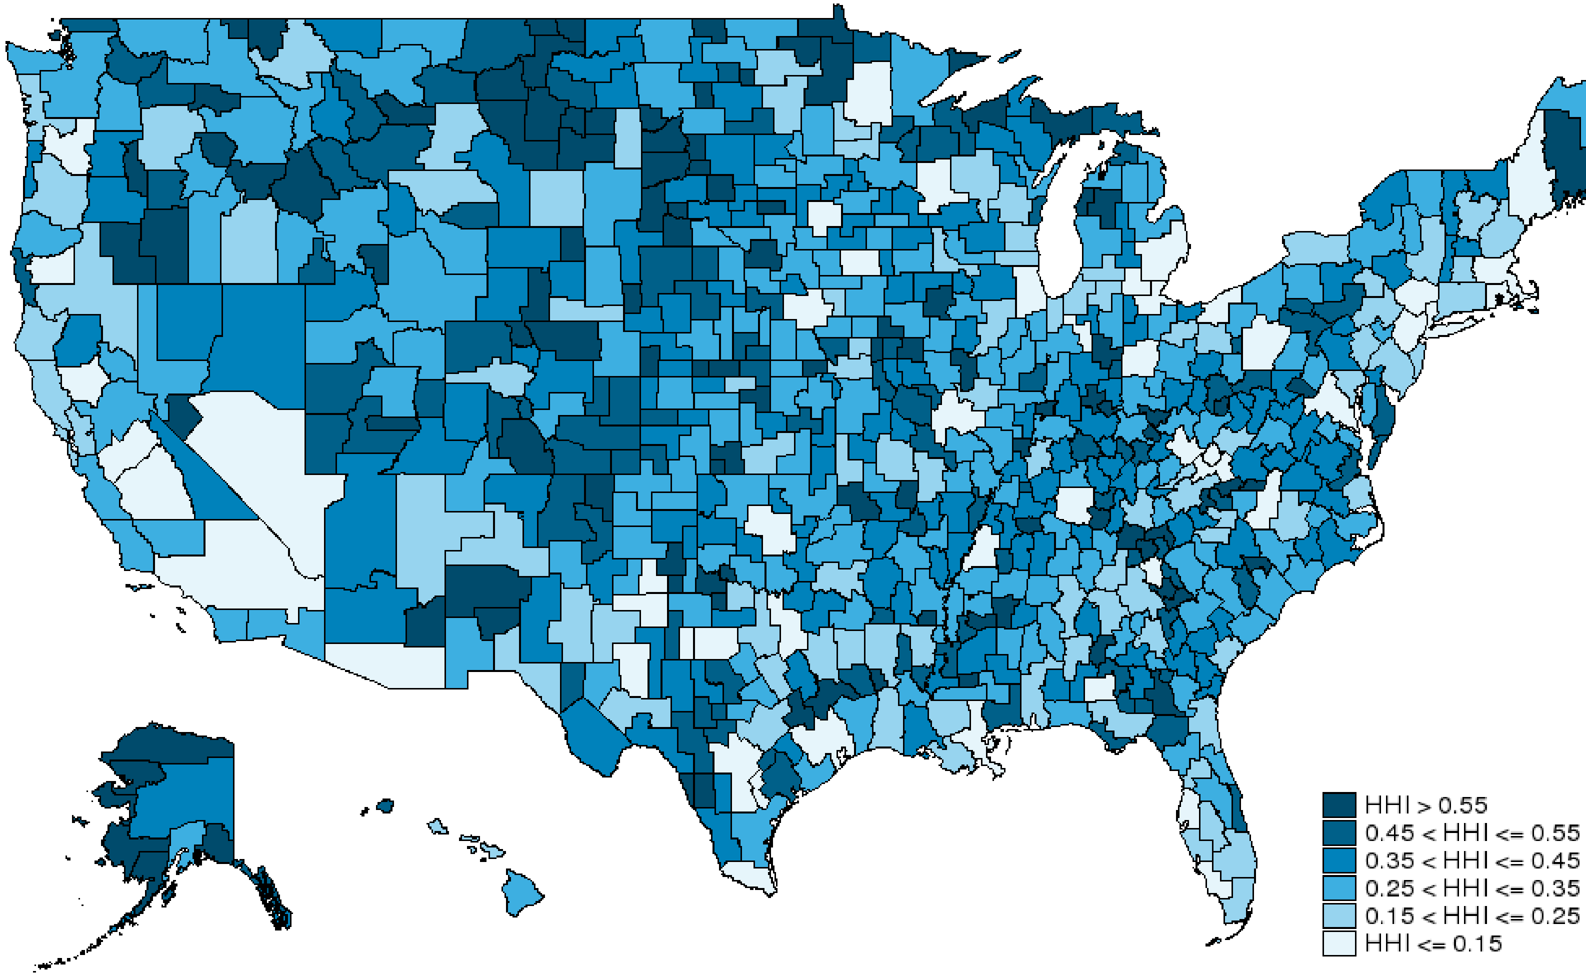

Source: Longitudinal Business Database, 1976

Note: Map plots the average HHI across four-digit NAICS industries within each commuting zone in 1976. Each commuting zones has had random noise drawn from a Laplace distribution added to its true value before being categorized.

Figure 10: Average Concentration Across Industries within Commuting Zones, 2015

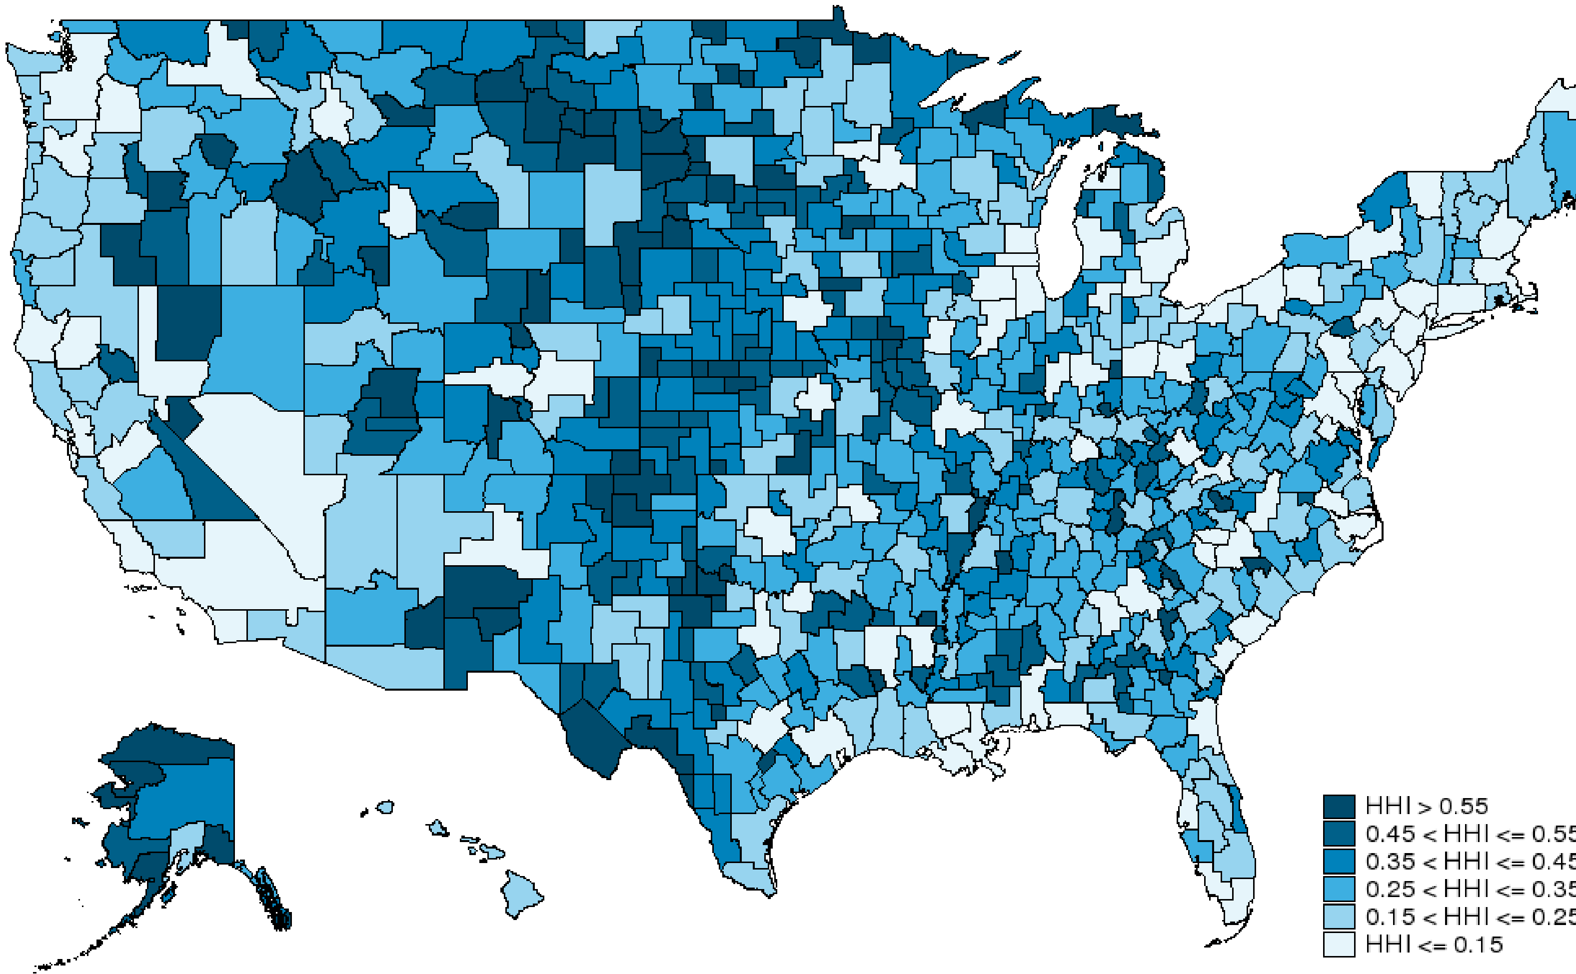

Source: Longitudinal Business Database, 2015

Note: Map plots the average HHI across four-digit NAICS industries within each commuting zone in 2015. Each commuting zones has had random noise drawn from a Laplace distribution added to its true value before being categorized.

Figure 11: Changes in the Log of Average Concentration Across Industries within Commuting Zones, 1976–1990

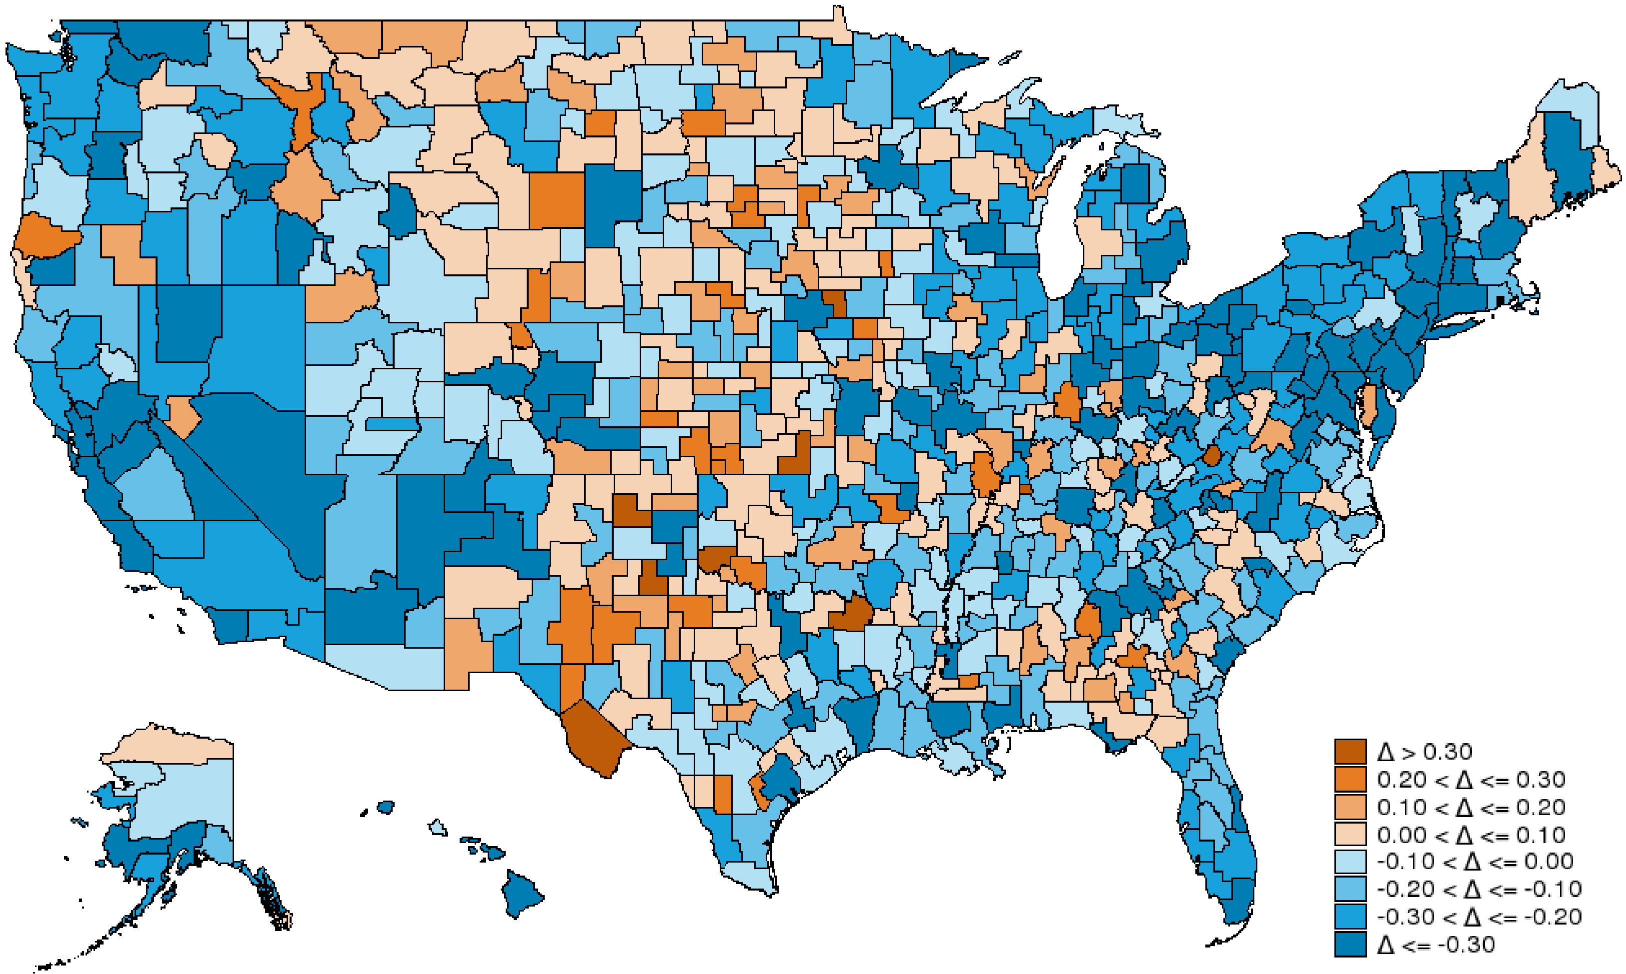

Source: Longitudinal Business Database, 1976 and 1990
Note: Map plots the change in the average HHI (represented by Δ in the legend) across four-digit NAICS industries within each commuting zone between 1976 and 1990. Each commuting zones has had random noise drawn from a Laplace distribution added to its true value before being categorized.

Figure 12: Changes in the Log of Average Concentration Across Industries within Commuting Zones, 1990–2005

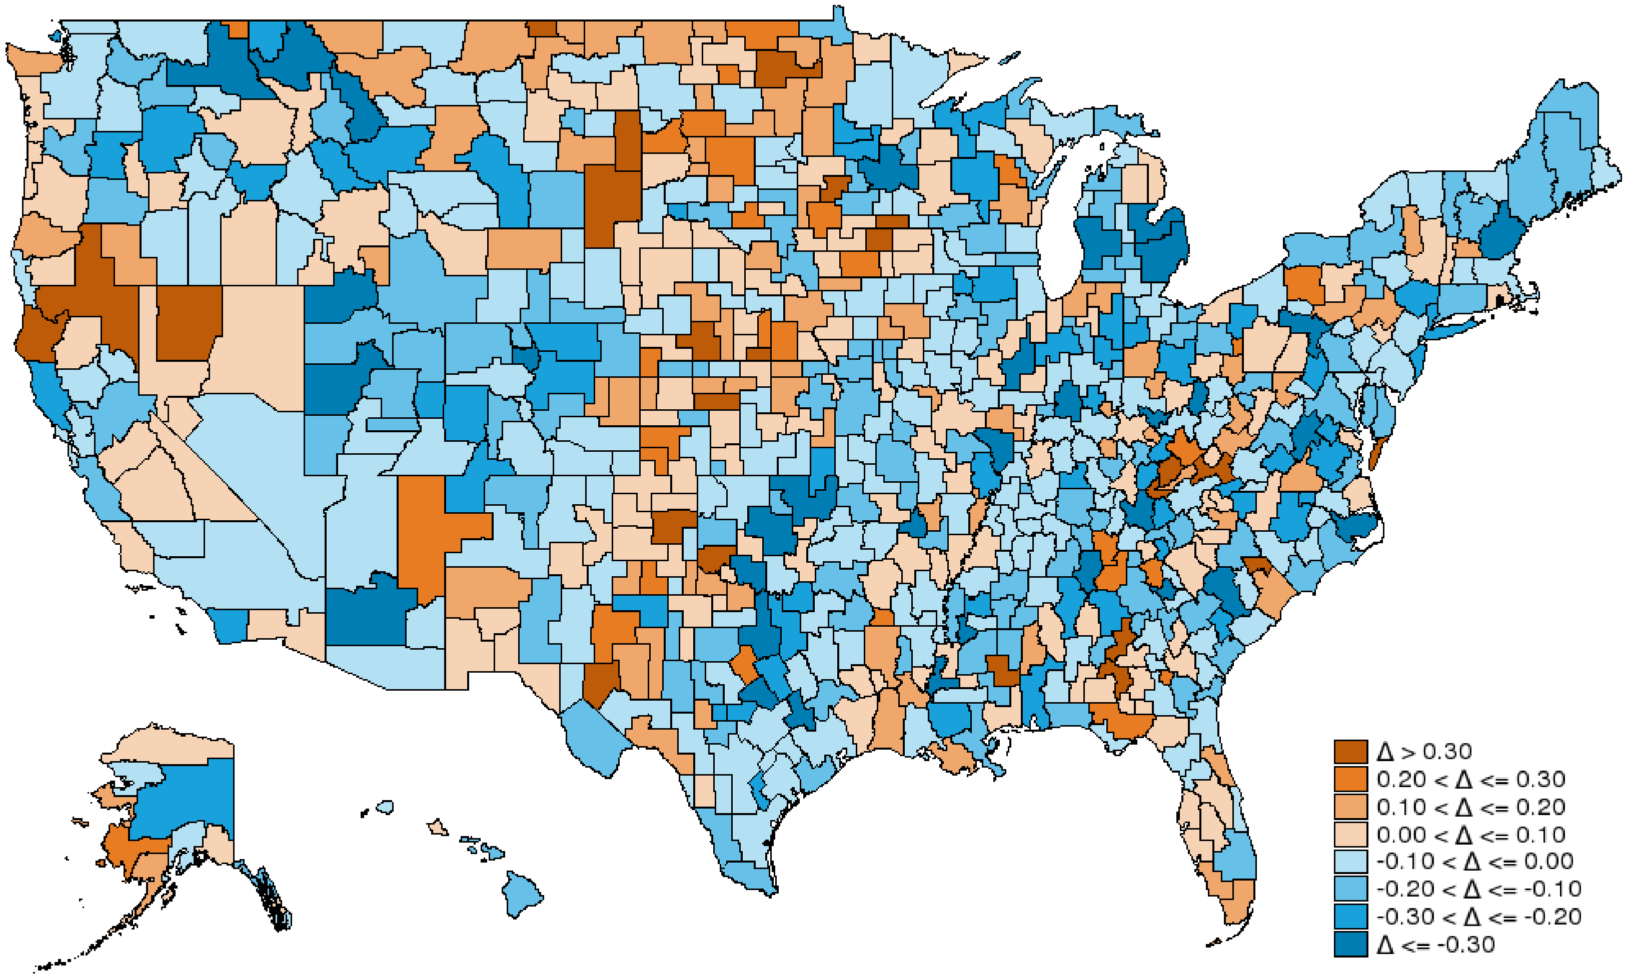

Source: Longitudinal Business Database, 1990 and 2005
Note: Map plots the change in the average HHI (represented by Δ in the legend) across four-digit NAICS industries within each commuting zone between 1990 and 2005. Each commuting zones has had random noise drawn from a Laplace distribution added to its true value before being categorized.

Figure 13: Changes in the Log of Average Concentration Across Industries within Commuting Zones, 2005–2015

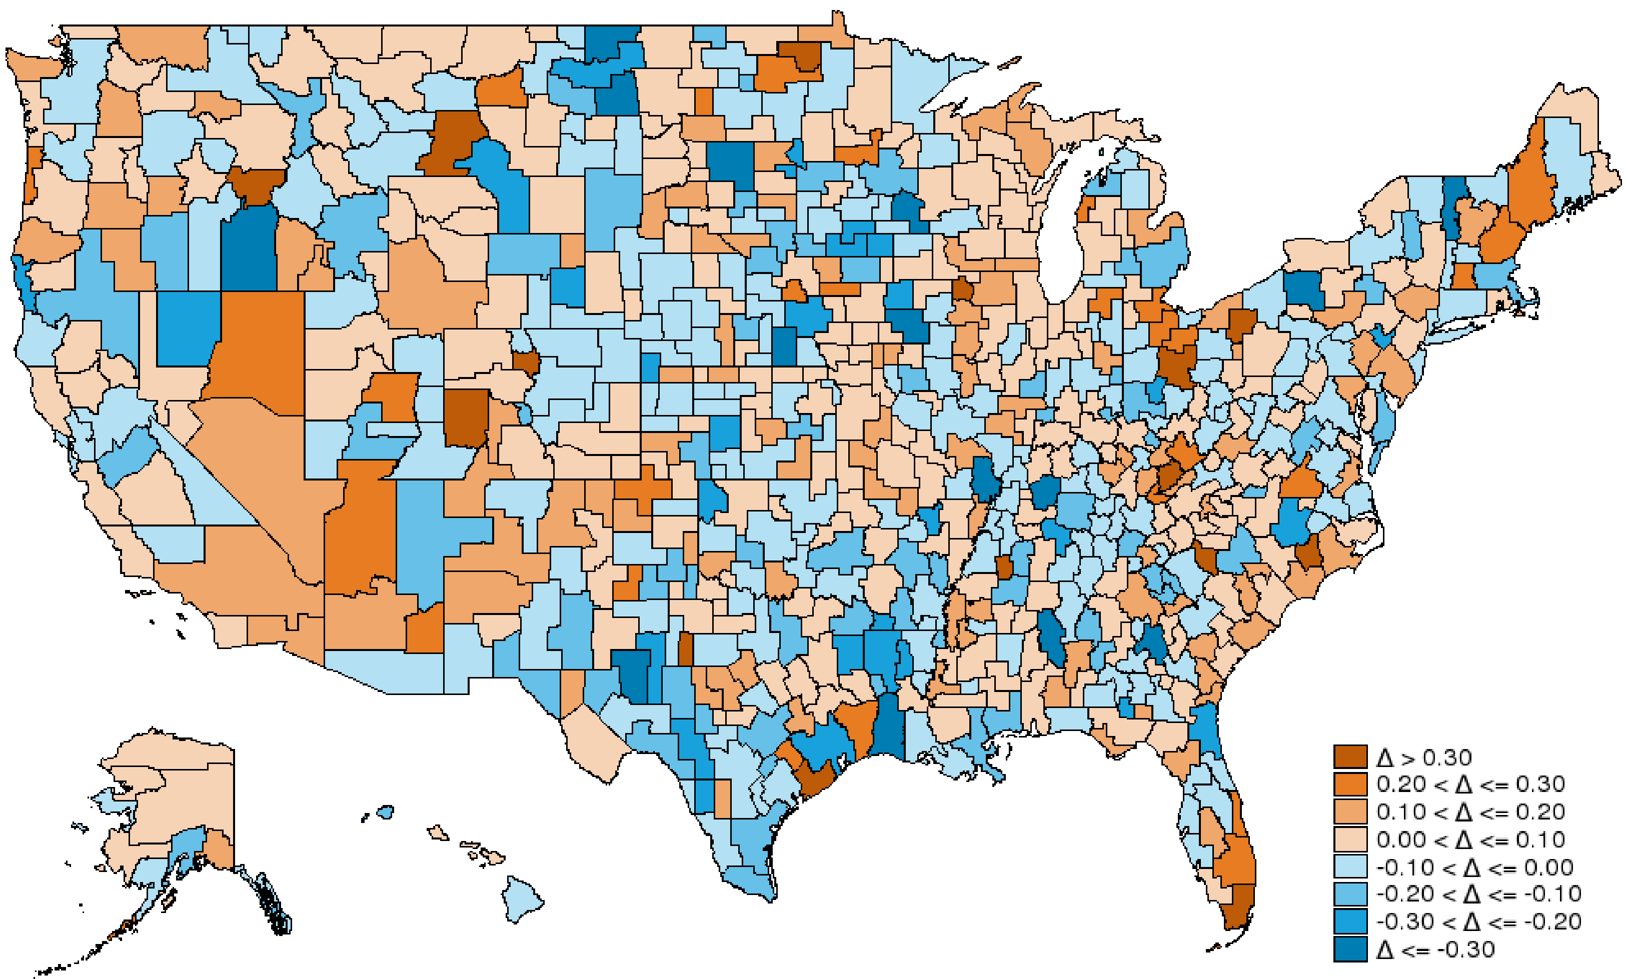

Source: Longitudinal Business Database, 2005 and 2015
Note: Map plots the change in the average HHI (represented by Δ in the legend) across four-digit NAICS industries within each commuting zone between 2005 and 2015. Each commuting zones has had random noise drawn from a Laplace distribution added to its true value before being categorized.

Figure 14: Changes in the Log of Average Concentration Across Industries within Commuting Zones, 1976–2015

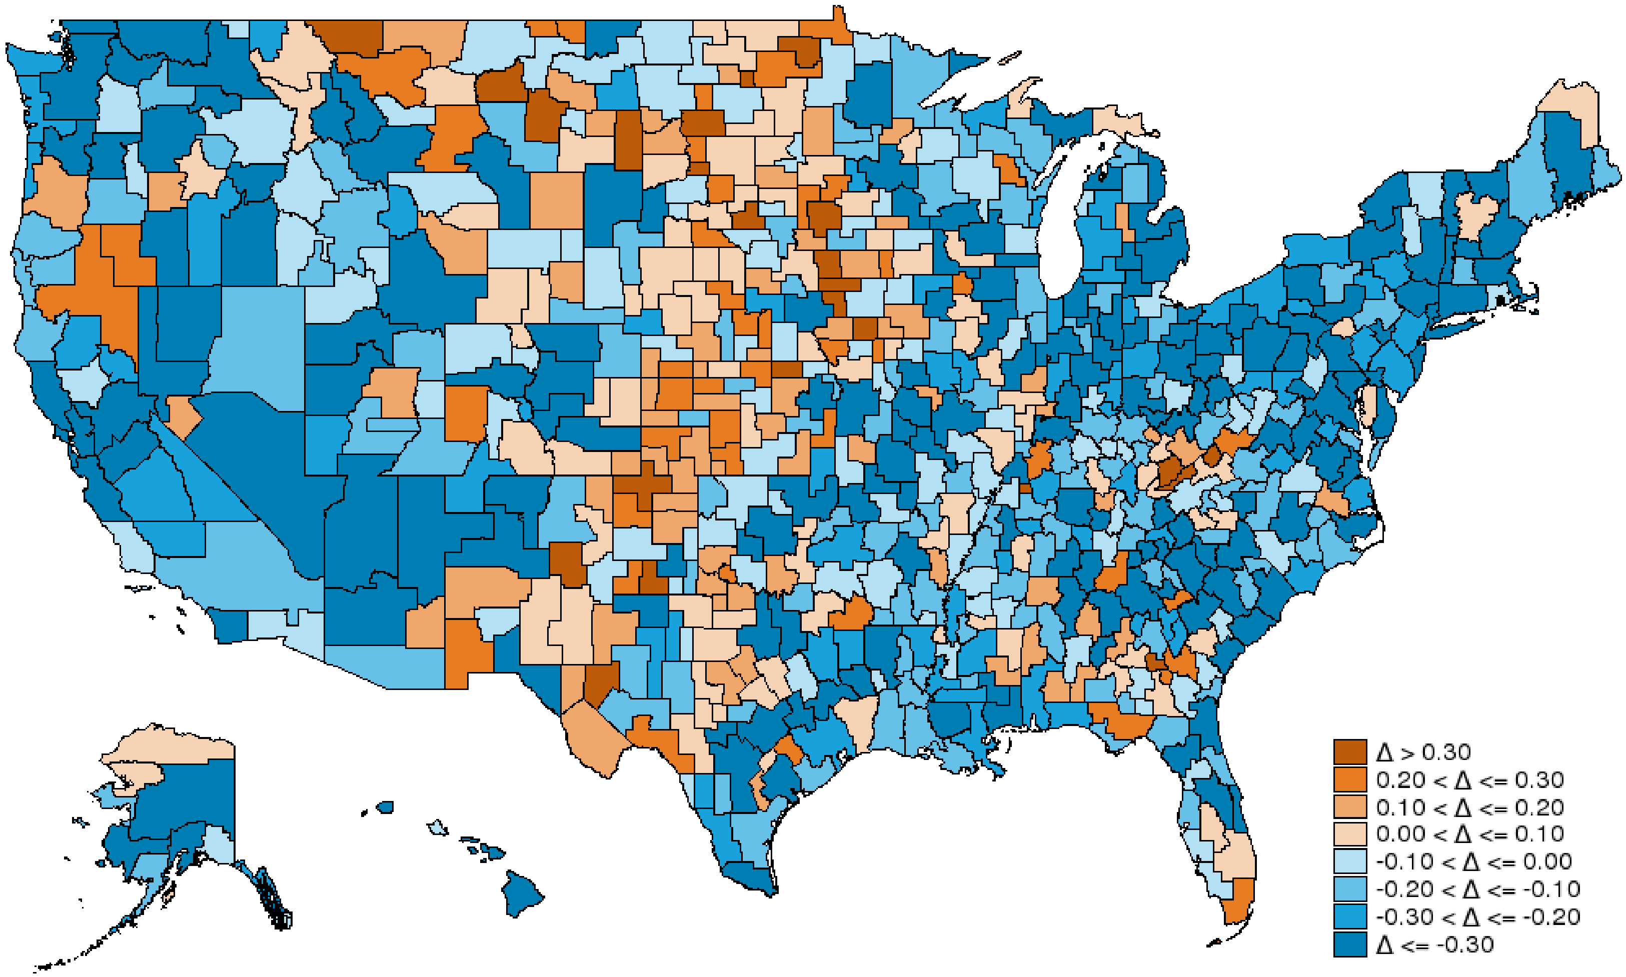

Source: Longitudinal Business Database, 1976 and 2015
Note: Map plots the change in the average HHI (represented by Δ in the legend) across four-digit NAICS industries within each commuting zone between 1976 and 2015. Each commuting zones has had random noise drawn from a Laplace distribution with parameter $\varepsilon = 15$ added to its true value before being categorized.

Figure 15: Changes in Mean Earnings versus Changes in Log Mean Industrial Concentration

(a) LBD Earnings, CZ Level, 1976–2015

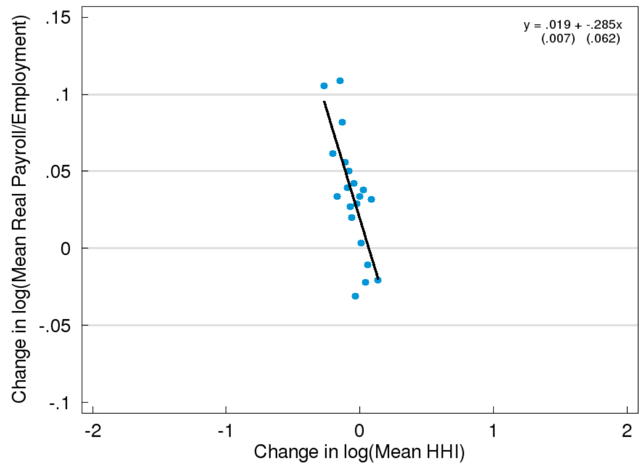

(b) LBD Earnings, CZ-Industry Level, 1976–2015

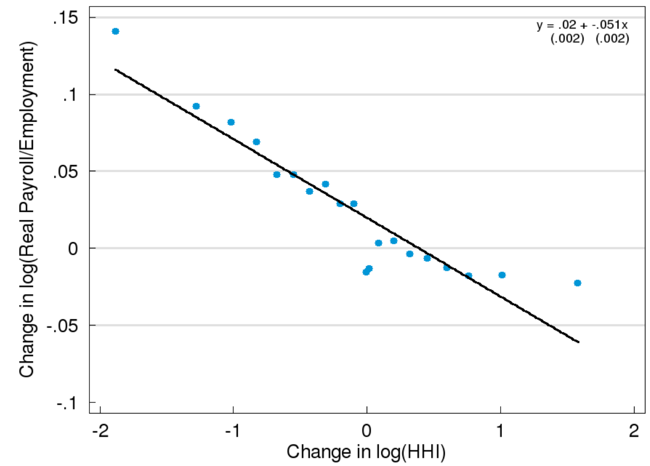

(c) LBD Earnings, CZ-Industry Level, 2005–2015

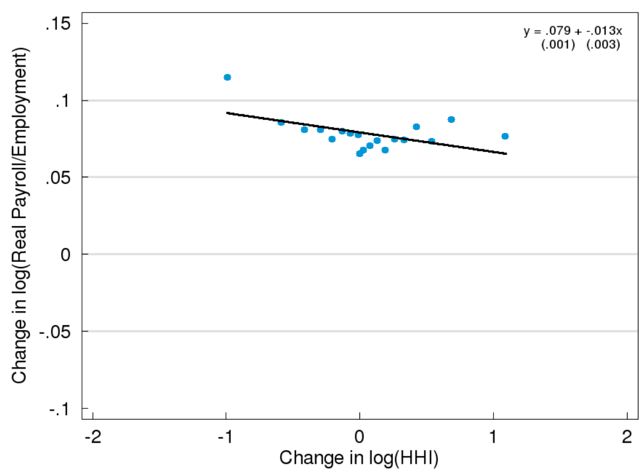

(d) W-2 Earnings, CZ-Industry Level, 2005–2015

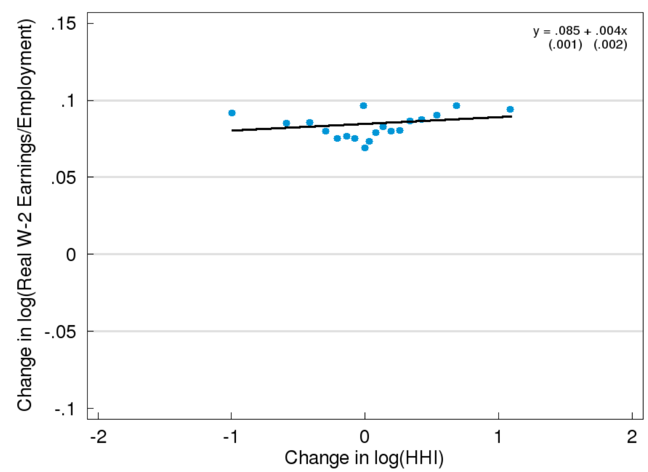

Source: Longitudinal Business Database and Form W-2, 1976, 2005, and 2015
Note: Figures plot changes in mean earnings against changes in local industrial concentration between the indicated years. Changes are calculated at the indicated level and then aggregated into twenty equal-sized bins, divided according the values of the change in industrial concentration. Earnings are obtained from the LBD in panels (a), (b), and (c), and from Form W-2 in panel (d).

Figure 16: Effects of Industrial Concentration on Key Percentiles of the Earnings Distribution

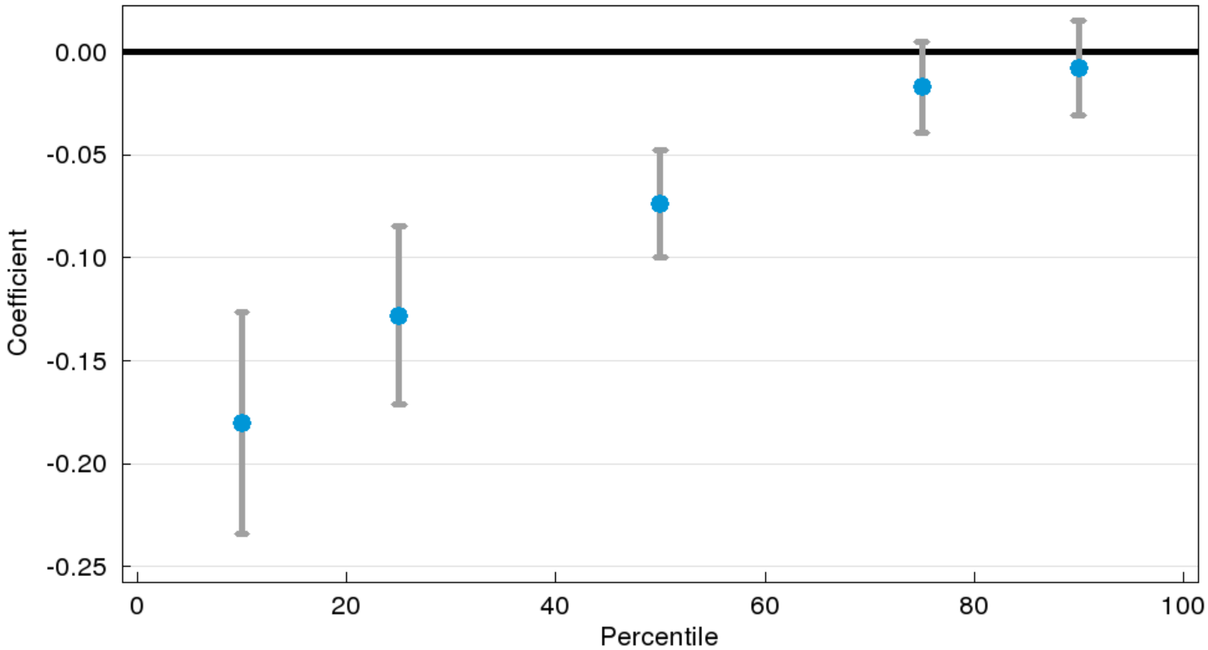

Source: Longitudinal Business Database and Form W-2, 2005–2015
Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the log of the values of key percentiles of the earnings distribution within markets on the log of local industrial concentration as measured by the HHI. Regressions include market and commuting zone by year fixed effects. Regressions are employment-weighted. Coefficients represent elasticities.

Figure 17: Effects of Industrial Concentration on Earnings, by Demographic Group

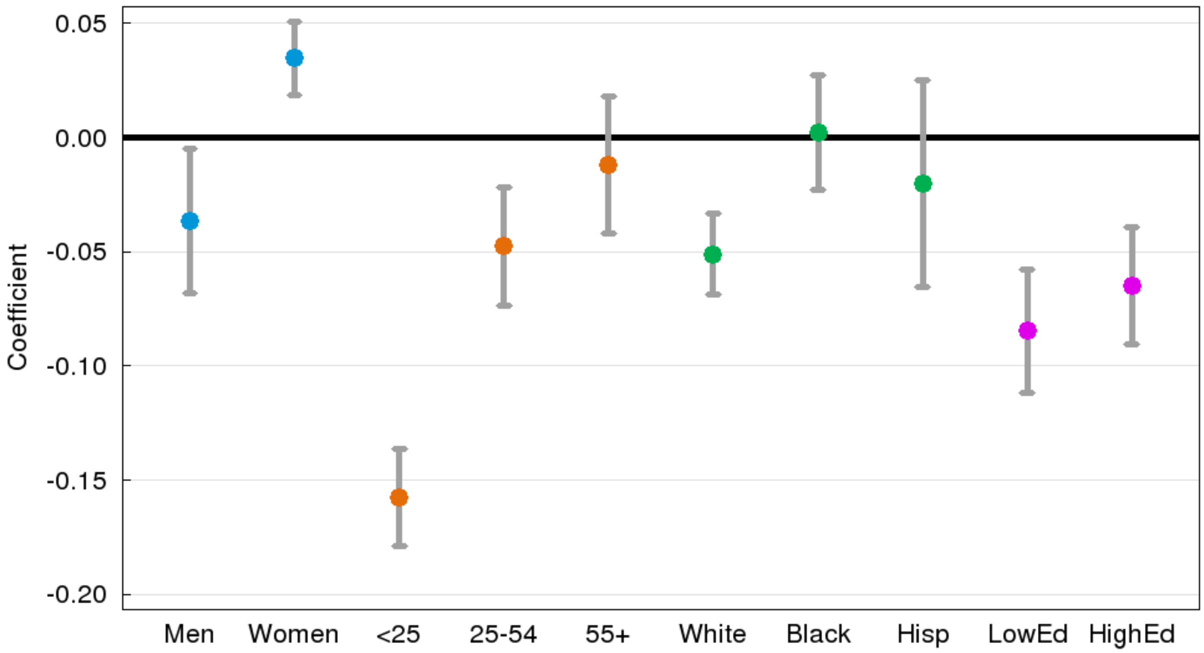

Source: Longitudinal Business Database, Form W-2, and American Community Survey, 2005 through 2015; Decennial Census, 2000 and 2010; Census Numident. For more information on the American Community Survey, see census.gov/acs.

Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the log of mean earnings within markets on the log of local industrial concentration as measured by the HHI for demographic groups identified on the x-axis. Regressions include market and commuting zone by year fixed effects. Regressions are employment-weighted. Coefficients represent elasticities. The White and Black categories refer to non-Hispanic White and non-Hispanic Black. The "Hisp" category includes Hispanics of any race. The "LowEd" category includes individuals with a high school diploma or less, while the "HighEd" category includes individuals who have at least attended some college.

Figure 18: Effects of Industrial Concentration on the 90/10 Earnings Ratio, by Demographic Group

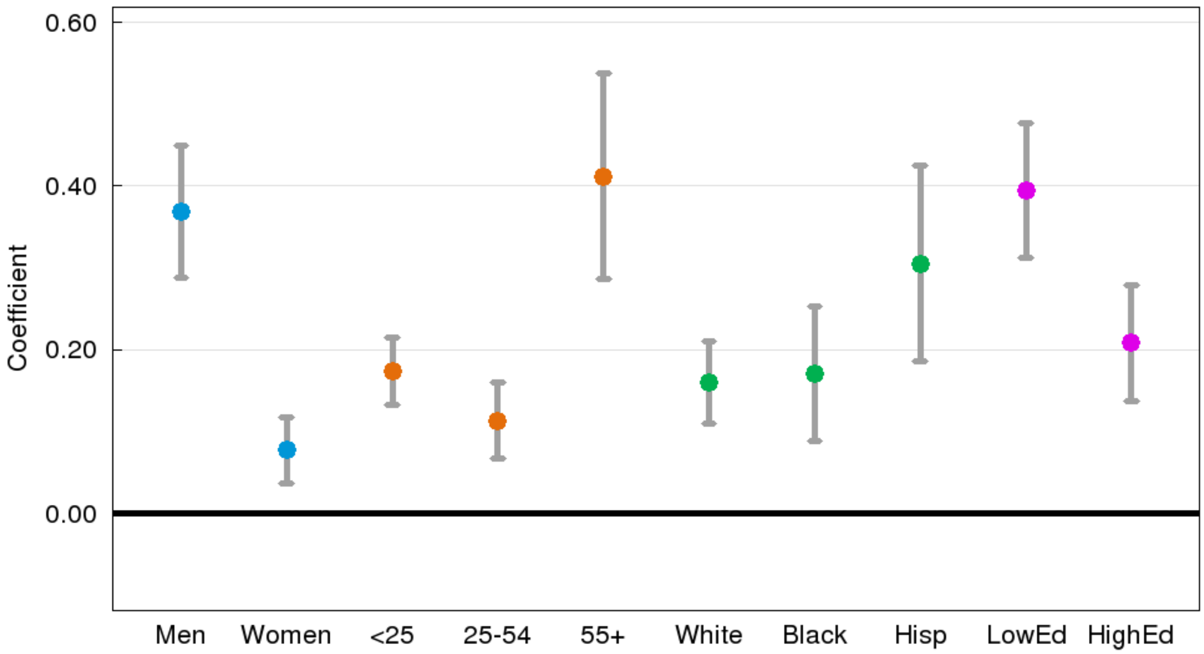

Source: Longitudinal Business Database, Form W-2, and American Community Survey, 2005 through 2015; Decennial Census, 2000 and 2010; Census Numident. For more information on the American Community Survey, see census.gov/acs.
Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the log of ratio of the 90th percentile of earnings to the 10th percentile of earnings within markets on the log of local industrial concentration as measured by the HHI for demographic groups identified on the x-axis. Regressions include market and commuting zone by year fixed effects. Regressions are employment-weighted. Coefficients represent elasticities. The White and Black categories refer to non-Hispanic White and non-Hispanic Black. The "Hisp" category includes Hispanics of any race. The "LowEd" category includes individuals with a high school diploma or less, while the "HighEd" category includes individuals who have at least attended some college.

Figure 19: Effects of Industrial Concentration on Earnings Percentile Ratios, by Demographic Group

(a) 50/10 Ratio

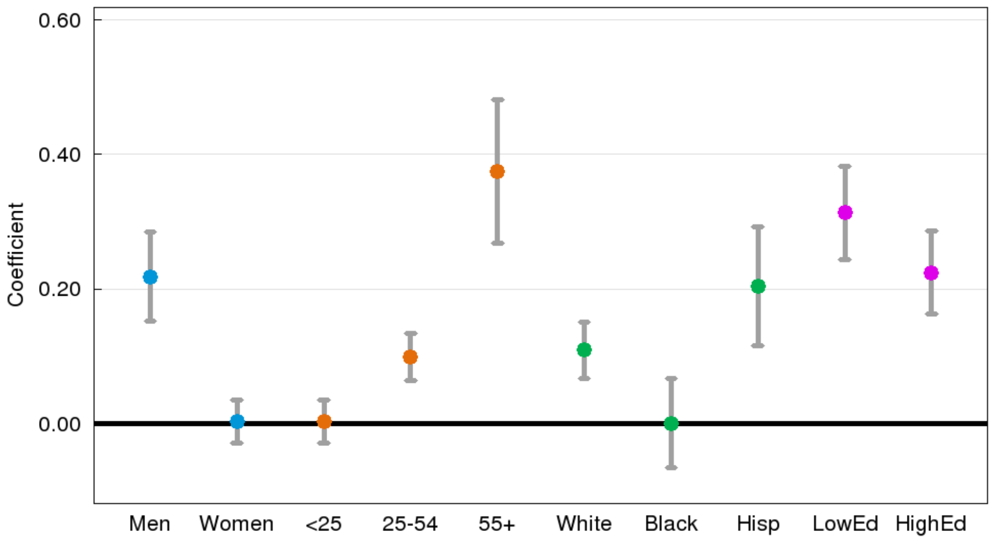

(b) 90/50 Ratio

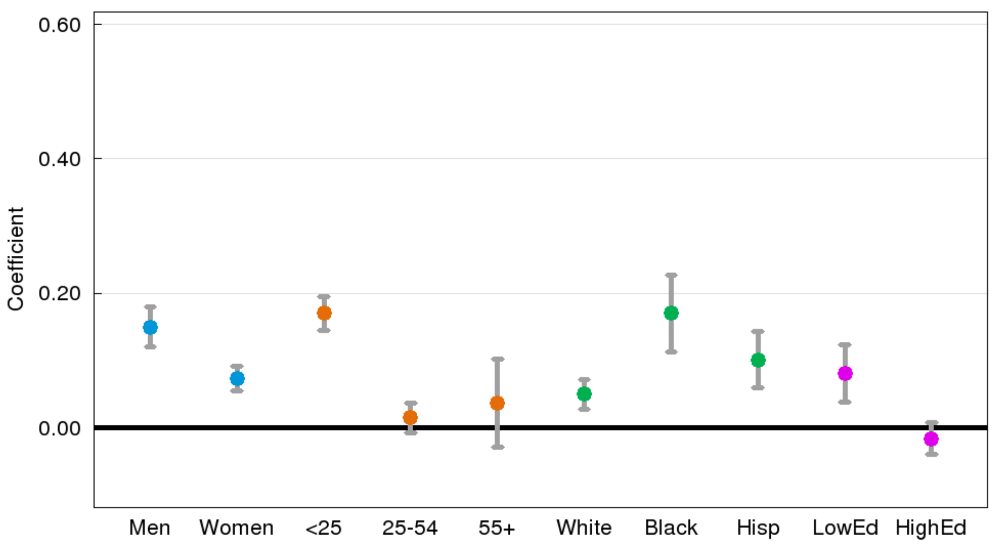

Source: Longitudinal Business Database, Form W-2, and American Community Survey, 2005 through 2015; Decennial Census, 2000 and 2010; Census Numident. For more information on the American Community Survey, see census.gov/acs.
Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the log of ratio of the indicated percentiles of the earnings distribution within markets on the log of local industrial concentration as measured by the HHI for demographic groups identified on the x-axis. Regressions include market and commuting zone by year fixed effects. Regressions are employment-weighted. Coefficients represent elasticities. The White and Black categories refer to non-Hispanic White and non-Hispanic Black. The "Hisp" category includes Hispanics of any race. The "LowEd" category includes individuals with a high school diploma or less, while the "HighEd" category includes individuals who have at least attended some college.

Figure 20: Effects of Industrial Concentration on the Gini Coefficient, by Demographic Group

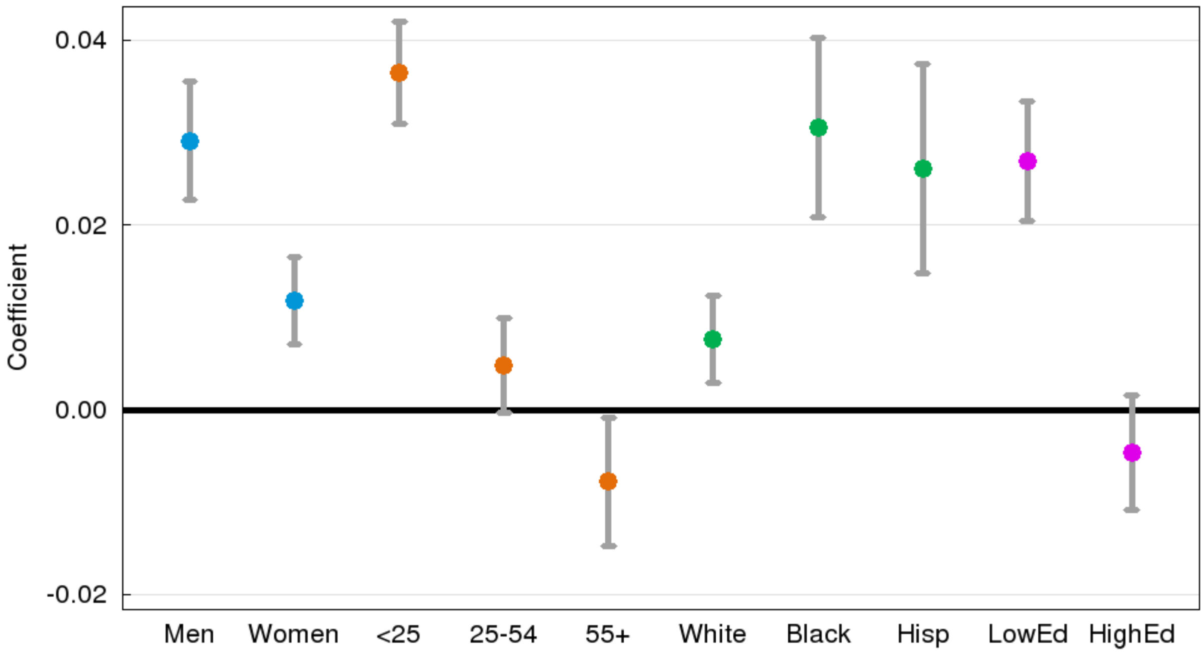

Source: Longitudinal Business Database, Form W-2, and American Community Survey, 2005 through 2015; Decennial Census, 2000 and 2010; Census Numident. For more information on the American Community Survey, see census.gov/acs.
Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the log of the Gini coefficient within markets on the log of local industrial concentration as measured by the HHI for demographic groups identified on the x-axis. Regressions include market and commuting zone by year fixed effects. Regressions are employment-weighted. Coefficients represent semi-elasticities. The White and Black categories refer to non-Hispanic White and non-Hispanic Black. The "Hisp" category includes Hispanics of any race. The "LowEd" category includes individuals with a high school diploma or less, while the "HighEd" category includes individuals who have at least attended some college.

Figure 21: Effects of Industrial Concentration on Relative Earnings Mobility, by Length of Horizon

(a) Baseline Specification

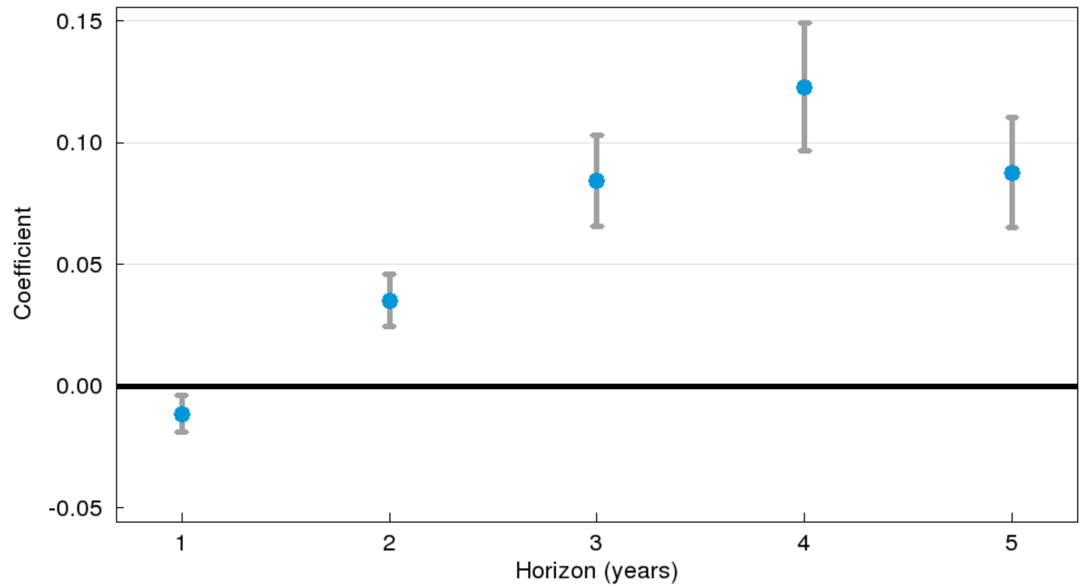

(b) With Market Trends

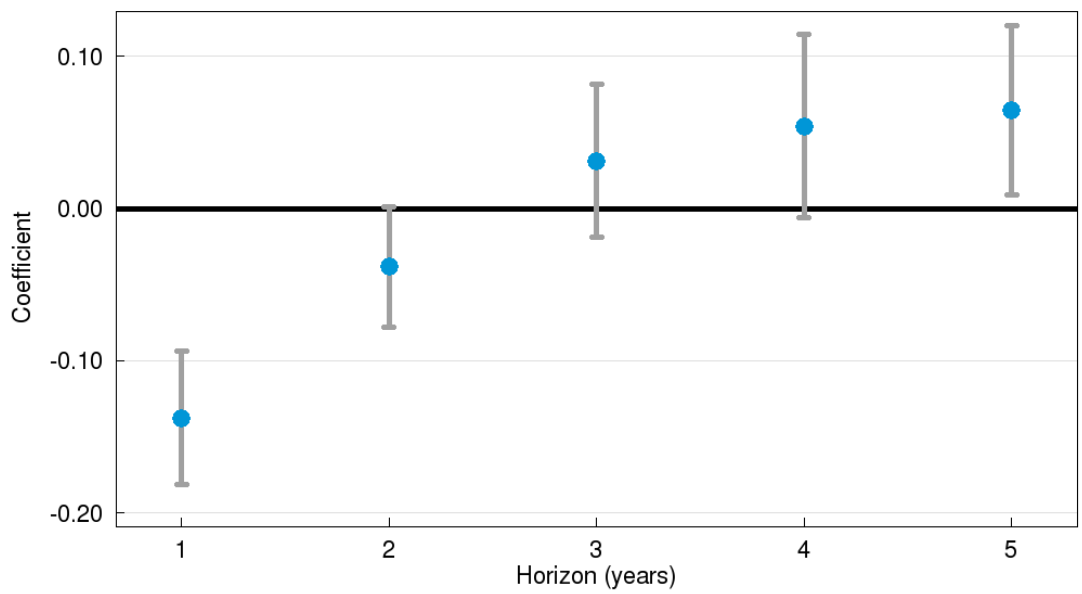

Source: Longitudinal Business Database and Form W-2, 2005–2015
Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the rank-rank coefficient of W-2 earnings estimated within markets over the horizon indicated on the x-axis on the log of local industrial concentration as measured by the HHI. Regressions include market and commuting zone by year fixed effects. Regressions in panel (b) also include market-specific linear trends. Regressions are employment weighted. Coefficients represent semi-elasticities.

Figure 22: Effects of Industrial Concentration on Absolute Earnings Mobility, by Length of Horizon

(a) Baseline Specification

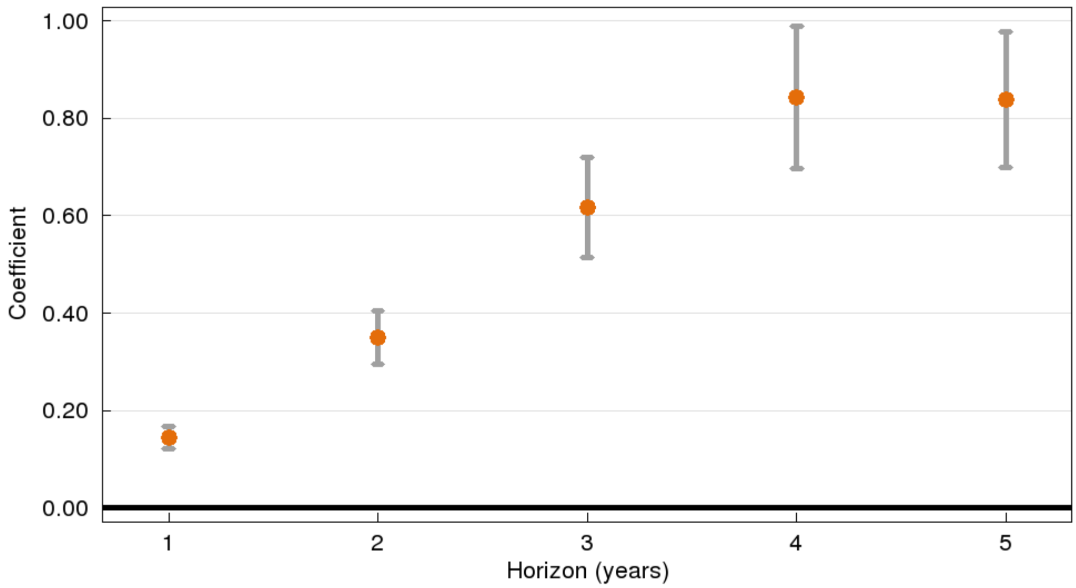

(b) With Market Trends

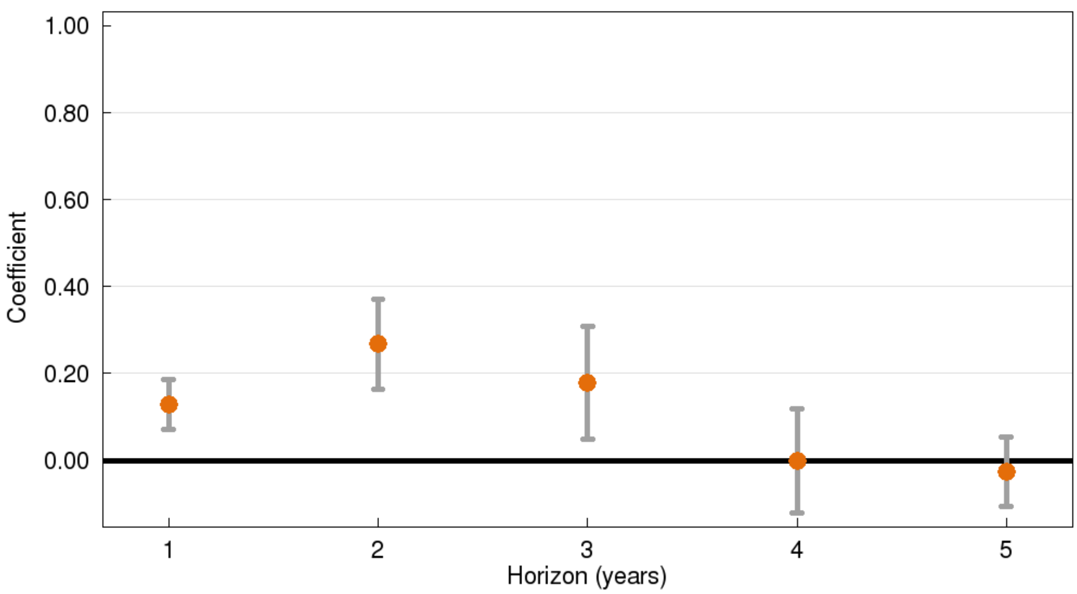

Source: Longitudinal Business Database and Form W-2, 2005–2015

Note: Figure plots regression coefficients and 95 percent confidence intervals from mean regressions of the change in log mean earnings within markets over the horizon indicated on the x-axis on the log of local industrial concentration as measured by the HHI. Regressions include market and commuting zone by year fixed effects. Regressions in panel (b) also include market-specific linear trends. Regressions are employment-weighted. Coefficients represent semi-elasticities.

Tables

Table 1: Effects of Industrial Concentration on Earnings, OLS Estimation

VARIABLES	(1)	(2)	(3)	(4)
log(HHI)	-0.108***	-0.0561***	0.00645***	0.00742***
	(0.00660)	(0.00368)	(0.00211)	(0.00117)
Observations	5,446,000	1,527,000	1,519,000	1,519,000
R-squared	0.658	0.972	0.983	0.872
Years	76-15	05-15	05-15	05-15
Earnings Measure	LBD	LBD	W-2	W-2
Weighted	Yes	Yes	Yes	No
Market FEs	Yes	Yes	Yes	Yes
CZ by Year FEs	Yes	Yes	Yes	Yes

Source: Longitudinal Business Database, 1976–2015; Form W-2, 2005–2015
Note: Table reports OLS regression estimates of the effect of local industrial concentration, as measured by the HHI, on log mean earnings. Earnings measures are constructed using either employment and payroll data from the LBD or earnings data from Form W-2, as indicated. Columns represent separate regressions, which include the indicated years of data and fixed effects. Regressions are employment-weighted as indicated. Coefficients represent elasticities. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 2: First Stage Regressions, 1976–2015

VARIABLES	(1)	(2)	(3)	(4)	(5)
$log(HHI^{-m})$	1.064***	0.748***	0.829***	0.827***	0.466***
	(0.0120)	(0.0201)	(0.0174)	(0.0173)	(0.0166)
Observations	5,450,000	5,450,000	5,446,000	5,446,000	5,446,000
R-squared	0.504	0.773	0.930	0.932	0.956
Year FEs	No	Yes	Yes	No	No
CZ FEs	No	Yes	No	No	No
Industry FEs	No	Yes	No	No	No
Market FEs	No	No	Yes	Yes	Yes
CZ by Year FEs	No	No	No	Yes	Yes
Market Trends	No	No	No	No	Yes
F-stat	7824	1389	2265	2284	791

Source: Longitudinal Business Database, 1976–2015
Note: Table reports regression estimates of the relationship between local industrial concentration, as measured by the HHI, and its instrument, the leave-one-out mean of the HHI across other markets in the same industry. Columns represent separate regressions, which include the indicated fixed effects in addition to the instrument. Regressions are employment-weighted. Coefficients represent elasticities. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 3: First Stage Regressions, 2005–2015

VARIABLES	(1)	(2)	(3)	(4)	(5)
$log(HHI^{-m})$	1.062***	-0.328***	0.503***	0.505***	0.192***
	(0.0130)	(0.0786)	(0.0303)	(0.0300)	(0.0226)
Observations	1,531,000	1,531,000	1,527,000	1,527,000	1,527,000
R-squared	0.537	0.792	0.974	0.974	0.985
Year FEs	No	Yes	Yes	No	No
CZ FEs	No	Yes	No	No	No
Industry FEs	No	Yes	No	No	No
Market FEs	No	No	Yes	Yes	Yes
CZ by Year FEs	No	No	No	Yes	Yes
Market Trends	No	No	No	No	Yes
F-stat	6667	17	276	284	73

Source: Longitudinal Business Database, 2005–2015
Note: Table reports regression estimates of the relationship between local industrial concentration, as measured by the HHI, and its instrument, the leave-one-out mean of the HHI across other markets in the same industry. Columns represent separate regressions, which include the indicated fixed effects in addition to the instrument. Regressions are employment-weighted. Coefficients represent elasticities. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 4: First Stage Regressions, 2005–2015, Markets with W-2 Earnings

VARIABLES	(1)	(2)	(3)	(4)	(5)
$log(HHI^{-m})$	1.053***	-0.131**	0.505***	0.505***	0.187***
	(0.0128)	(0.0640)	(0.0280)	(0.0274)	(0.0204)
Observations	1,522,000	1,522,000	1,519,000	1,519,000	1,519,000
R-squared	0.540	0.801	0.975	0.975	0.986
Year FEs	No	Yes	Yes	No	No
CZ FEs	No	Yes	No	No	No
Industry FEs	No	Yes	No	No	No
Market FEs	No	No	Yes	Yes	Yes
CZ by Year FEs	No	No	No	Yes	Yes
Market Trends	No	No	No	No	Yes
F-stat	6747	4	326	339	84

Source: Longitudinal Business Database and Form W-2, 2005–2015
Note: Table reports regression estimates of the relationship between local industrial concentration, as measured by the HHI, and its instrument, the leave-one-out mean of the HHI across other markets in the same industry. Columns represent separate regressions, which include the indicated fixed effects in addition to the instrument. Regressions are employment-weighted. Coefficients represent elasticities. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 5: Effects of Industrial Concentration on Mean Earnings

VARIABLES	(1)	(2)	(3)	(4)
log(HHI)	-0.0512**	-0.00857	-0.0324***	-0.109***
	(0.0200)	(0.0122)	(0.0117)	(0.0121)
Observations	5,446,000	1,527,000	1,519,000	1,519,000
R-squared	0.657	0.972	0.983	0.871
Years	76-15	05-15	05-15	05-15
Earnings Measure	LBD	LBD	W-2	W-2
Weighted	Yes	Yes	Yes	No
Market FEs	Yes	Yes	Yes	Yes
CZ by Year FEs	Yes	Yes	Yes	Yes

Source: Longitudinal Business Database, 1976–2015; Form W-2, 2005–2015
Note: Table reports instrumental variables regression estimates of the effect of local industrial concentration, as measured by the HHI, on log mean earnings. Earnings measures are constructed using either employment and payroll data from the LBD or earnings data from Form W-2, as indicated. Columns represent separate regressions, which include the indicated years of data and fixed effects. Regressions are employment-weighted as indicated. Coefficients represent elasticities. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 6: Effects of Industrial Concentration on Earnings Inequality

VARIABLES	(1) 90/10	(2) 50/10	(3) 90/50	(4) Gini
log(HHI)	0.173***	0.107***	0.0659***	0.0124***
	(0.0265)	(0.0210)	(0.0123)	(0.00273)
Observations	1,519,000	1,519,000	1,519,000	1,519,000
R-squared	0.895	0.841	0.880	0.940
Market FEs	Yes	Yes	Yes	Yes
CZ by Year FEs	Yes	Yes	Yes	Yes

Source: Longitudinal Business Database and Form W-2, 2005–2015
Note: Table reports instrumental variables regression estimates of the effect of local industrial concentration, as measured by the HHI, on measures of earnings inequality, constructed using earnings data from Form W-2. The dependent variables are the logs of the ratios of the 90th and 10th (Column 1), 50th and 10th (Column 2), or 90th and 50th (Column 3) percentiles of the earnings distribution, and the Gini coefficient (Column 4). Columns represent separate regressions, which include the indicated years of data and fixed effects. Regressions are employment-weighted as indicated. Coefficients in columns 1-3 represent elasticities, while the coefficient in column 4 is a semi-elasticity. Sample sizes and statistic values have been rounded for disclosure avoidance.

Table 7: Effects of Industrial Concentration on Earnings Outcomes, Combined Non-Tradable and Construction Sector

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	HHI	Mean Earnings	90/10	90/50	50/10	Gini
$log(HHI^{-m})$	0.344***					
	(0.0285)					
log(HHI)		-0.184***	0.396***	0.0976***	0.298***	0.0148***
		(0.0278)	(0.0691)	(0.0223)	(0.0538)	(0.00506)
Observations	333,000	333,000	333,000	333,000	333,000	333,000
R-squared	0.976	0.970	0.867	0.936	0.767	0.933
Market FEs	Yes	Yes	Yes	Yes	Yes	Yes
CZ by Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	145.0					

Source: Longitudinal Business Database and Form W-2, 2005–2015

Note: Table reports instrumental variables regression estimates of the effect of local industrial concentration, as measured by the HHI, on measures of earnings and inequality, constructed using earnings data from Form W-2, within the combined non-tradable and construction sector, as defined by Mian and Sufi (2014). The first column reports the first-stage regression. In the subsequent columns, the dependent variables are the log of mean earnings (Column 2), the logs of the ratios of the 90th and 10th (Column 3), 50th and 10th (Column 4), or 90th and 50th (Column 5) percentiles of the earnings distribution, and the Gini coefficient (Column 6). Columns represent separate regressions, which include the indicated years of data and fixed effects. Regressions are employment-weighted. Coefficients in columns 2-5 represent elasticities, while the coefficient in column 6 is a semi-elasticity. Sample sizes and statistic values have been rounded for disclosure avoidance.