

RICHARD T. ELY LECTURE

Work of the Past, Work of the Future[†]

By DAVID H. AUTOR*

One of the enduring paradoxes that has accompanied the rise of wage inequality over the last four decades in industrialized economies is the sustained *fall* in real wages experienced by less-educated workers.¹ This fall is seen in the United States, Germany, and the United Kingdom, among other nations (Dustmann, Ludsteck, and Schönberg 2009; Acemoglu and Autor 2011; Blundell et al. 2018). It is illustrated for the United States in Figure 1, which plots cumulative changes in real log weekly wage and salary earnings of full-time, full-year workers between 1963 and 2017.

The progression of inequality over these five and a half decades can be roughly divided into three epochs: (i) the ten-year interval between 1963 and 1972, when real wages rose robustly and evenly among all education by gender groups; (ii) the interregnum between 1973 and

1979 when, following the first US oil shock, real earnings growth stagnated throughout the distribution; and (iii) the era of secularly rising wage inequality from 1980 forward, where wages rose robustly among the most-educated and fell in real terms among the least-educated—most strikingly, among men with less than a bachelor’s degree.² Among high school dropouts, high school graduates, and those with some college, real weekly earnings among full-time male workers in 2018 were 10 to 20 log points below their real levels in 1980. While the evolution of real wages was less adverse among non-college women, there was a 15-plus year period between 1981 and 1997 when women with high school or lower education earned less than their counterparts in 1980.

To understand why the fall in real earnings of less-educated workers is paradoxical, consider the rapid increase in the educational attainment of the US labor force during the time period documented in Figure 2. Among both sexes, the share of labor hours supplied by workers with high school or lower education fell from more than 75 percent in 1963 to less than 40 percent in 2017.³ Conversely, the share of labor hours supplied by workers with a bachelor’s or post-college degree rose from less than 15 percent to more than 35 percent.⁴ The rapid increase

*Department of Economics, MIT, Morris and Sophie Chang Building, 50 Memorial Drive, E52-438, Cambridge, MA 02142, and NBER (email: dautor@mit.edu). I am indebted to Daron Acemoglu, David Dorn, Juliette Fournier, Claudia Goldin, Colin Gray, Gordon Hanson, Lawrence Katz, James Poterba, Anna Salomons, and Evan Soltas for ideas, insights, and critiques that enriched this work. Anne Beck, Emiel Van Bezooijen, Pepe (Jose Ignacio Velarde) Morales, Edwin Song, and Sunny (Liang) Tan provided abundant and ingenious hard work to put these ideas to the test. I thank Accenture LLP, the IBM Global Universities Program, the Schmidt Futures Foundation, and the Smith Richardson Foundation, for generous financial support.

[†]Go to <https://doi.org/10.1257/pandp.20191110> to visit the article page for additional materials and author disclosure statement(s).

¹While inequality has many facets—inequality of labor versus non-labor income, income versus consumption inequality, transient versus permanent earnings inequality, “between-group” versus residual inequality, inequality of top incomes (i.e., “the 1 percent”) versus inequality elsewhere in the distribution (“the other 99 percent”)—my focus here is on the inequality of wage and salary earnings by education and sex. For discussion of other aspects of inequality, see Lemieux (2006); Autor (2014); Piketty and Saez (2014); and Aguiar and Bils (2015).

²Within this post-1979 epoch, two subperiods of robust wage growth deserve mention: the years 1995 to 2000, corresponding to the so-called “dot-com” boom; and the sustained economic expansion from 2012 to present, following the end of the Great Recession. While the 1995–2000 interval saw rapid wage growth without any reduction in wage inequality, Figure 1 suggests that during the current expansion, wages are rising particularly rapidly among the least educated.

³Note that these hours shares should be understood as an equilibrium outcome, not as a measure of latent labor supply.

⁴Also visible is that the rise in educational attainment—particularly college degree holding—has been far greater among women than men. In 2017, 36 percent of male hours

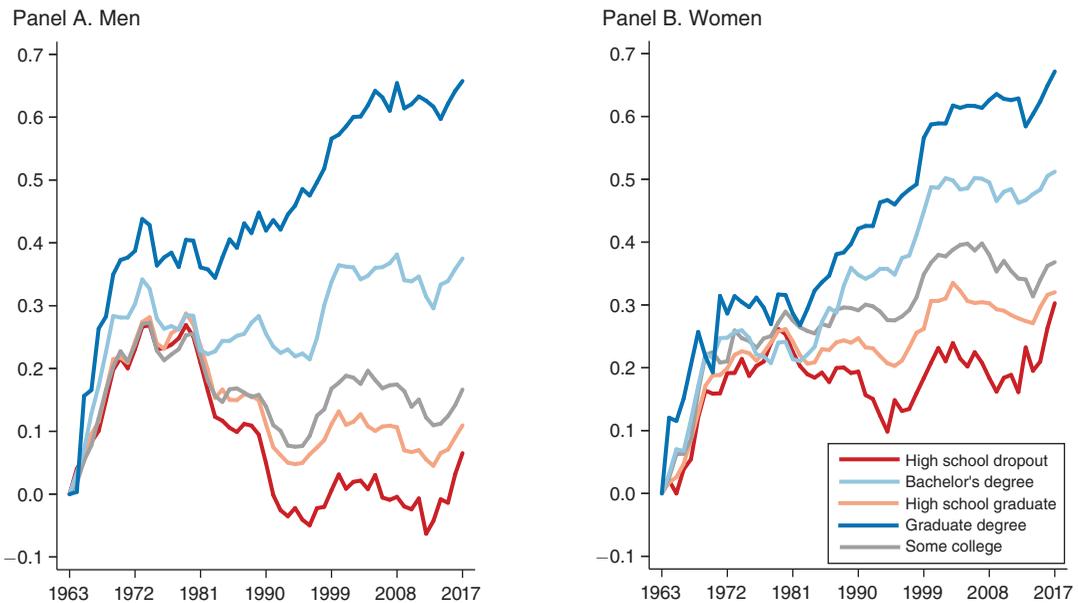


FIGURE 1. CUMULATIVE CHANGE IN REAL WEEKLY EARNINGS OF WORKING-AGE ADULTS AGES 18–64, 1963–2017

Notes: Figure uses March Current Population Survey Annual Social and Economic Supplement data for earnings years 1963 to 2017. Series correspond to (composition-adjusted) mean log wage for each group, using data on full-time, full-year workers ages 16 to 64. The data are sorted into sex-education-experience groups of two sexes, five education categories (high school dropout, high school graduate, some college, college graduate, and post-college degree), and four potential experience categories (0–9, 10–19, 20–29, and 30–39 years). Educational categories are harmonized following the procedures in Autor, Katz, and Kearney (2008). Log weekly wages of full-time, full-year workers are regressed in each year separately by sex on dummy variables for four education categories, a quartic in experience, three region dummies, black and other race dummies, and interactions of the experience quartic with three broad education categories (high school graduate, some college, and college plus). The (composition-adjusted) mean log wage for each of the 40 groups in a given year is the predicted log wage from these regressions evaluated for whites, living in the mean geographic region, at the relevant experience level (5, 15, 25, or 35 years depending on the experience group). Mean log wages for broader groups in each year represent weighted averages of the relevant (composition-adjusted) cell means using a fixed set of weights, equal to the mean share of total hours worked by each group over 1963–2005. All earnings numbers are deflated by the chain-weighted (implicit) price deflator for personal consumption expenditures. Earnings of less than \$67/week in 1982 dollars (\$112/week in 2000 dollars) are dropped. Allocated earnings observations are excluded in earnings years 1967 forward using either family earnings allocation flags (1967–1974) or individual earnings allocation flags (1975 earnings year forward).

in overall educational attainment coupled with *rising* wages for college-educated workers is *not* hard to rationalize in a market-clearing supply and demand framework. Assuming that college and non-college workers are gross complements (i.e., their elasticity of substitution in production exceeds one), the simultaneous rise in the supply *and* wages of college-educated workers is consistent with ongoing skill-biased technological change that causes the productivity of—and hence demand for—college workers to grow

(even) more rapidly than the supply (see Katz and Murphy 1992; Acemoglu 2002; Goldin and Katz 2008).

It is far *harder* to rationalize the falling real wages of non-college workers in this same framework, however. If college and non-college workers are gross complements, as we have just assumed, an increase in the relative supply of college workers or a rise in their productivity should boost the productivity—and hence the wages—of non-college workers. This assertion is no more mysterious than the notion that capital complements labor, which implies that capital deepening should raise labor productivity and hence wages. But this outcome has

and 43 percent of female hours were supplied by workers with a bachelor's degree or greater.

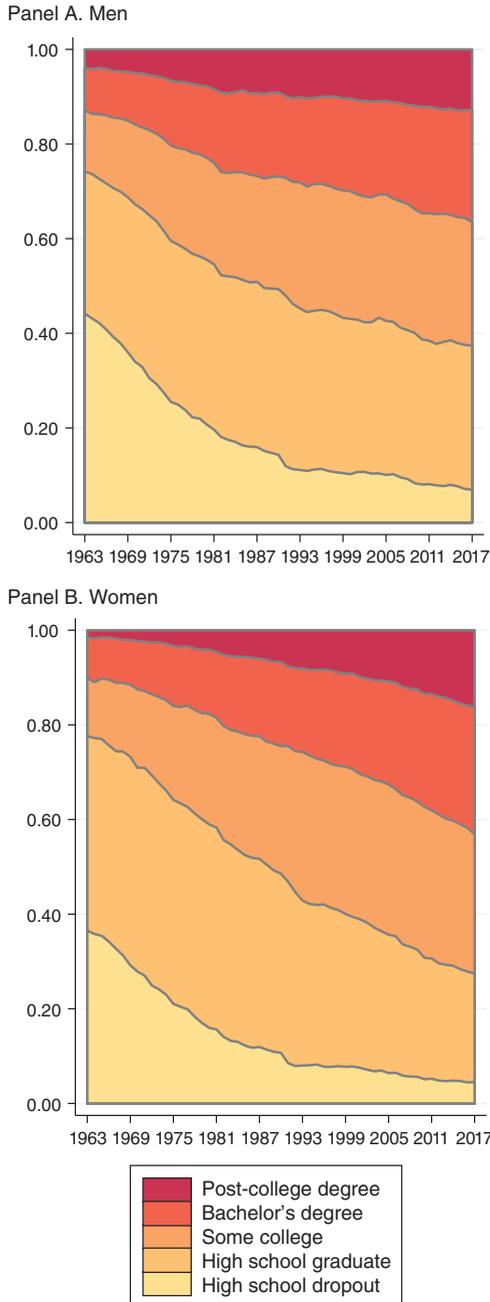


FIGURE 2. SHARE OF HOURS WORKED IN THE US ECONOMY BY EDUCATION GROUP, 1963–2017

Notes: Data source is as in Figure 1. Sample consists of all persons aged 16 to 64 who reported having worked at least one week in the earnings years, excluding those in the military. For each individual, hours worked are the product of usual hours worked per week and the number of weeks worked last year. Individual hours worked are aggregated using CPS sampling weights.

demonstrably not occurred. Over the course of nearly four decades, non-college workers in the United States have *not* as a group benefited from the rising supply and productivity of college-educated workers.

That paradox is the subject of this paper. There are of course *many* potential explanations, including but not limited to eroding union penetration and bargaining power, falling federal and state minimum wages, rising trade pressure accompanying China's rise as a manufacturing power, and the "fissuring" of the workplace, wherein less-educated workers no longer share in the gains from rising productivity and profitability in the core activities of their employers (Card 1996; Lee 1999; Autor, Dorn, and Hanson 2013, 2016; Weil 2014). Literature has shown that all of these factors are important. My focus here, however, is on one economic force that stands at the intersection of technological progress and worker productivity: occupational change. It is well understood that the structure of work in industrialized countries has polarized, with employment increasingly concentrated in high-education, high-wage occupations and low-education, low-wage occupations, at the expense of traditionally middle-skill career jobs (Autor, Katz, and Kearney 2006, 2008; Goos and Manning 2007; Autor 2013; Michaels et al. 2014; Goos, Manning, and Salomons 2014). Less widely recognized is the tight connection between this polarization and the changing structure of work and wages across *geographic regions*.⁵ While the polarization of the occupational structure has unfolded smoothly over four decades, it has *not* unfolded evenly across places. In the decades following WWII, there was a steep, positive urban gradient in the skill level and wage level of *non-college* jobs: Non-college urban adults disproportionately held middle-skill, blue-collar production and white-collar office, administrative, and clerical jobs. Because these workers labored in close collaboration with the high-skill, urban

⁵It has long been understood that cities and skills are deeply entwined (Glaeser and Mare 2001, Florida 2002). And to be sure, I am not the first to study differential polarization across places (cf. Autor 2013; Mazzolari and Ragusa 2013; Akerman, Gaarder, and Mogstad 2015). The primary contribution of this paper is to demonstrate the centrality of geography to both the prevalence of middle-skill jobs in earlier decades and to their steep decline in recent decades.

professional, managerial, and technical workers who oversaw factories and offices, middle-skill jobs for non-college workers were prevalent in cities and metropolitan areas but scarce in suburbs and rural labor markets.⁶

Starting in the 1970s, however, this distinctive feature of dense labor markets began to recede. As rising automation and international trade encroached on employment in production, administrative support, and clerical work, the non-college urban occupational skill gradient diminished and ultimately disappeared. While workers are vastly more educated and jobs are vastly more skill-intensive in today's cities compared to five decades ago, non-college workers in these places perform substantially *less* skilled work than they did decades earlier. Polarization thus reflects an *unwinding* of the distinctive structure of work in dense cities and metro areas relative to suburban and rural areas.

I argue here that this unwinding—concretely, the differential polarization of urban labor markets—has contributed profoundly to the decline of non-college wages documented in Figure 1. I sketch three mechanisms by which polarization may have contributed to falling non-college wages: it has shunted non-college workers from middle-skill career occupations that reward specialized and differentiated skills into traditionally low-education occupations that demand primarily generic skills; it has disproportionately depressed middle-wage employment among non-college workers in urban labor markets, thus directly reducing average non-college wages and—to a startling degree—attenuating the urban non-college wage premium that prevailed in earlier decades; and it has created an excess supply of less-educated workers that serves to depress non-college wages across occupations and geographic areas.

The paper proceeds as follows: Section I describes the advance of occupational polarization over four decades, documents the disproportionate prevalence of this polarization among non-college workers, and presents initial evidence that this *aggregate* phenomenon can partly

⁶Of course non-college workers in both urban and non-urban labor markets performed traditionally low-education, low-wage manual labor, transportation, construction, and in-person service jobs. Distinctively, many non-college workers in urban labor markets held middle-skill jobs.

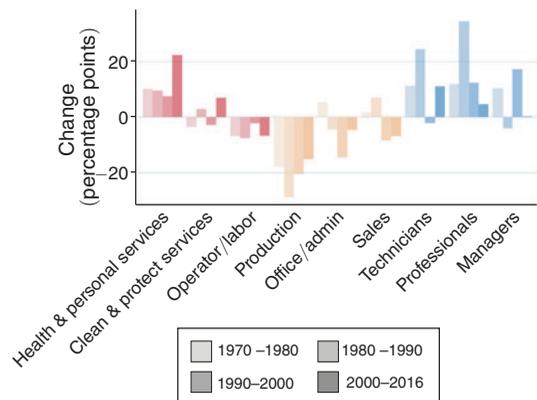


FIGURE 3. PERCENT CHANGES IN OCCUPATIONAL EMPLOYMENT SHARES AMONG WORKING-AGE ADULTS, 1970–2016

Notes: Data source is as in Figure 1. Sample consists of all persons aged 16 to 64 who reported having worked at least one week in the earnings years, excluding those in the military. For each individual, hours worked are the product of usual hours worked per week and the number of weeks worked last year. Individual hours worked are aggregated using CPS sampling weights. Occupational classifications are harmonized following Dorn (2009), and updated through 2017. Further details on occupational classification are provided in the supporting data supplements for the paper.

account for falling non-college wages. Section II reports a key new result: Occupational polarization has been disproportionately urban and reflects the undoing of a previously robust urban occupational skill gradient among non-college workers. Complementing this evidence, Section III documents the striking secular decline in the urban wage premium among non-college workers, and shows how the urban concentration of occupational polarization further contributes to the falling wages of non-college workers. Section IV concludes by asking whether the growth of new urban occupations will lead to a renaissance of urban middle-skill work.

I. Occupational Polarization

A. The Big Picture

Figure 3 depicts the familiar polarization of the occupational structure of the US labor market that has unfolded over the course of more than four decades. The nine exhaustive and mutually exclusive occupational categories depicted in

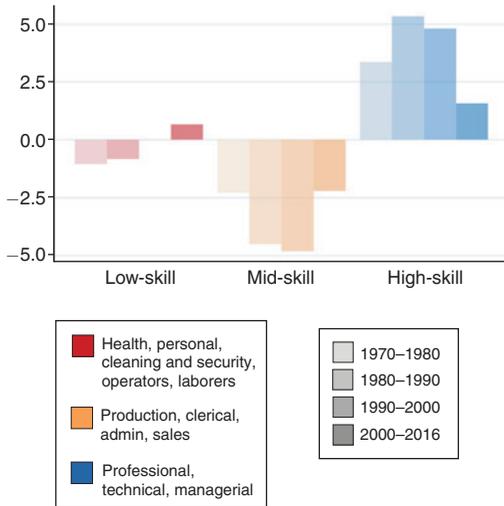


FIGURE 4. CHANGES IN OCCUPATIONAL EMPLOYMENT SHARES AMONG WORKING-AGE ADULTS, 1970–2016

Notes: Data source is as in Figure 1. Sample consists of all persons aged 16 to 64 who reported having worked at least one week in the earnings years, excluding those in the military. For each individual, hours worked are the product of usual hours worked per week and the number of weeks worked last year. Individual hours worked are aggregated using CPS sampling weights. Occupational classifications are harmonized following Dorn (2009), and updated through 2017. Further details on occupational classification are provided in the supporting data supplements for the paper.

this figure are ordered from lowest to highest by mean log wage level. These nine categories are further clustered into three broad occupational groups depicted in distinct colors: service and manual occupations; production, office, administrative, and sales occupations; and technical, professional, and managerial occupations. The “barbell” shape of this figure reflects the secular bifurcation of the occupational structure in the United States (and many other industrial economies) into high-education, high-wage professional, managerial, and technical occupations, on the one hand, and non-credentialed and typically low-paid service and laborer occupations on the other (see Autor, Katz, and Kearney 2006; Goos and Manning 2007; Goos, Manning, and Salomons 2014; Autor 2015; Acemoglu and Restrepo 2017; Alabdulkareem et al. 2018).⁷

⁷Plotted bars correspond to the proportional change in the share of employment in each category; smaller categories

Figure 4 brings these patterns into sharper focus by aggregating the nine occupation categories into three broad clusters of manual and service occupations (“low” skill); production, office, and sales occupations (“middle” skill); and professional, technical, and managerial occupations (“high” skill). At the start of this interval in 1970, US employment was roughly evenly divided among these three categories: 31.4 percent of total hours were in low-skill occupations, 38.4 percent were middle-skill, and 30.2 percent were high-skill.⁸ Over the subsequent four-and-a-half decades, middle-skill employment fell steeply, from 38.4 to 23.3 percent of hours. This trend might be concerning were it not the case that more than the entirety of this fall was offset by a rise in employment in high-skill occupations, which grew from 30.2 to 46.2 percent of hours. Meanwhile, the share of employment in low-skill occupations *fell* by almost a percentage point, from 31.4 to 30.6 percent. Thus, in aggregate, occupational polarization appears to be a case of the middle-class joining the upper-class, which is not something that economists should worry about.

Figure 5 tempers that conclusion. Among college workers (those with some college or higher education), occupational movement has been modestly, though not uniformly, upward. Between 1980 and 2016, the fraction of college workers in high-skill occupations rose from 57.2 percent to 60.7 percent, the share in middle-skill occupations fell from 27.1 to 20.2 percent, and the share in low-skill occupations increased from 15.6 to 19.0 percent.⁹ Thus, occupational polarization among college workers has broken roughly evenly between reallocation toward traditionally high- and low-skill jobs.

Among non-college workers (those with high school or lower education), the picture is radically different. In 1980, employment of non-college workers was roughly split between low- and middle-skill occupations, with 42 percent in the

can have large growth rates without accounting for a large change in employment and vice versa for larger categories.

⁸In 1980, these shares were 30.3, 36.1, and 33.5 percent, respectively.

⁹When splitting the data according to educational attainment, I focus on the period from 1980 forward because incompatibilities in Census occupational codes between 1970 and 1980 are greatly amplified when the data are split by education.

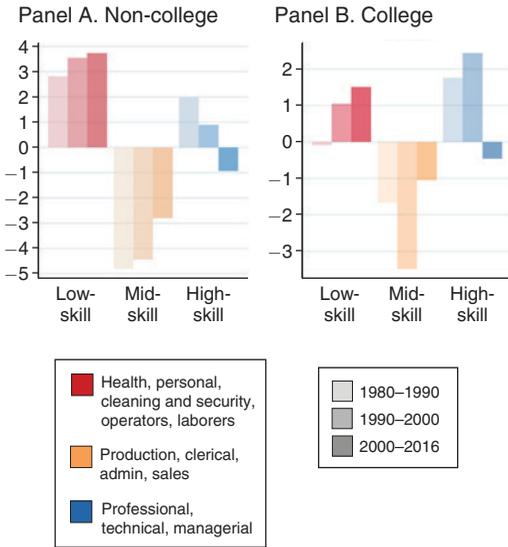


FIGURE 5. CHANGES IN OCCUPATIONAL EMPLOYMENT SHARES AMONG WORKING-AGE ADULTS, 1970–2016

Notes: Data source is as in Figure 1. Sample consists of all persons aged 16 to 64 who reported having worked at least one week in the earnings years, excluding those in the military. For each individual, hours worked are the product of usual hours worked per week and the number of weeks worked last year. Individual hours worked are aggregated using CPS sampling weights. Occupational classifications are harmonized following Dorn (2009), and updated through 2017. Further details on occupational classification are provided in the supporting data supplements for the paper.

former category, 43 percent in the latter, and the remaining 15 percent in high-skill occupations. Over the ensuing decades, the share of non-college employment in middle-skill occupations fell by 14 percentage points, from 43 to 29 percent, while the share in high-skill occupations very slightly increased, from 15.4 to 16.8 percent. Thus, the remaining 12.3 percentage points of decline in non-college middle-skill employment is explained by the movement of non-college workers from middle-skill into traditionally low-skill work. This is a qualitatively large change, and I will argue next that it has been economically consequential.¹⁰

¹⁰ Although somewhat counterintuitive, it is not a violation of adding up for the low-skill share to rise among both college and non-college workers and yet fall in aggregate. The resolution is that college workers are much less likely than non-college workers at all times to work in low-skill

A foundational assumption of the modern literature on skill demand, dating at least to Tinbergen (1974), is that technological progress complements—and hence raises demand for—educated workers. This framing might suggest that highly-educated workers should see their work transformed by technology. While this transformation has to some degree occurred, a clear takeaway from this descriptive analysis is that changes in the nature of work—many of which are technological in origin—have been far more profound and, arguably, far more disruptive for less-educated workers than they have been for more-educated workers. In broad strokes, the work performed by college adults has changed little over four decades. While they perform fewer middle-skill jobs than four decades earlier, this contraction has been modest, and it has been substantially offset by their upward movement in the occupational hierarchy. Among non-college workers, conversely, polarization has exerted pressure almost exclusively downward: Almost all occupational change among non-college workers reflects a movement from the middle toward the bottom of the occupational distribution. Thus, not only has technology change been transformational, it has been broadly deskilling—by which I mean that it has narrowed the set of jobs in which non-college workers perform specialized work that historically (and currently, as I show below) commanded higher pay levels.

B. A Simple Calibration

A natural question that follows from the data above is whether occupational change can in part explain the sharp, post-1980 divergence in real wages by education seen in Figure 1. To explore that question, I perform a simple partial-equilibrium calculation to construct counterfactual wage series in which I hold the occupational wage structure fixed at its 1978 level while allowing the distribution of workers by education and gender to shift across occupations as observed in the data.¹¹ Formally, we can write the change in the mean log wage of

occupations, and the fraction of college versus non-college workers is rising throughout this period.

¹¹ The precise interval of this calculation, 1978–2016, is determined by occupational comparability at the start of the interval and by the need to take centered means across three

education group j between two years t_0 and t_1 using the identity

$$\Delta \bar{w}_{j\tau} = \sum_k (\alpha_{jkt_1} \omega_{jkt_1} - \alpha_{jkt_0} \omega_{jkt_0}),$$

where the α_{jkt} terms correspond to the fraction of group j workers employed in occupation k in year t , while ω_{jkt} is their mean log wage in that occupation and year.¹² Manipulating this expression, I calculate a counterfactual wage series for each education group j that isolates the occupational change component (the $\Delta\alpha$ terms) while holding the wage structure (ω_{jkt_0}) terms fixed:

$$(1) \quad \Delta \tilde{w}_{j\tau} = \sum_k \bar{w}_{jkt_0} (\alpha_{jkt_1} - \alpha_{jkt_0}).$$

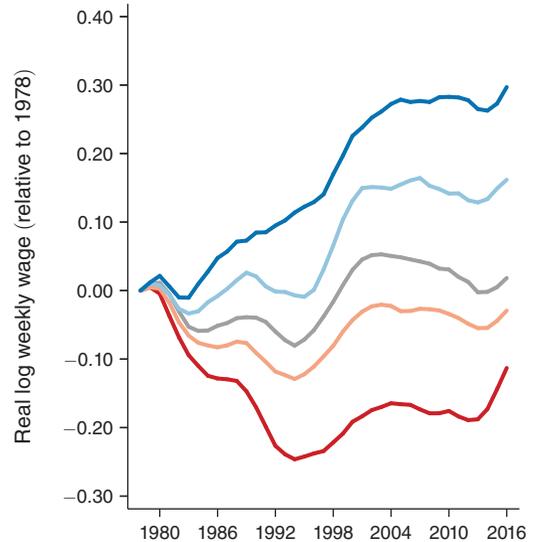
Figure 6 reports estimates of $\Delta \tilde{w}_{j\tau}$, alongside observed changes in mean log real hourly wages, by education group for earnings years 1978 through 2016. Superficially, this exercise appears to capture the general evolution of wages by education group. It clearly captures the fanning out of wage levels by education group, the fall in real wages of non-college workers, and also the notable divergence between earnings of college graduates and those with a post-college degree throughout this period—and particularly after the late 1990s. Closer inspection reveals two substantive disjunctures. A first is that the actual figure shows considerable real wage growth among workers with a college degree and higher, something that is absent from the counterfactual series. In part, this divergence is an artifact: since real wages used in the exercise are fixed at their 1978 levels, the counterfactual series omits any wage-augmenting productivity growth that occurs in the ensuing three decades.

The more serious departure between the series is evident from comparing the y-axis ranges of

years to smooth the series. The last year of data, 2017, is subsumed in the 2016 mean.

¹²The 12 occupational categories that comprise this calculation (and are used throughout the paper) are, from low to high mean wages: Agriculture and mining; Health services; Personal services; Cleaning and protective services; Construction and mechanics; Transportation; Production and operatives; Clerical and administrative support; Retail sales minus financial and advertising; Technicians, fire and police; Professionals plus financial and advertising sales; Managers and executives. In implementing this calculation, I further allow mean log wages and employment shares to differ by gender within education groups j .

Panel A. Actual wage changes between occupation reallocation + within occupation wage changes



Panel B. Counterfactual wage changes between occupation reallocation only

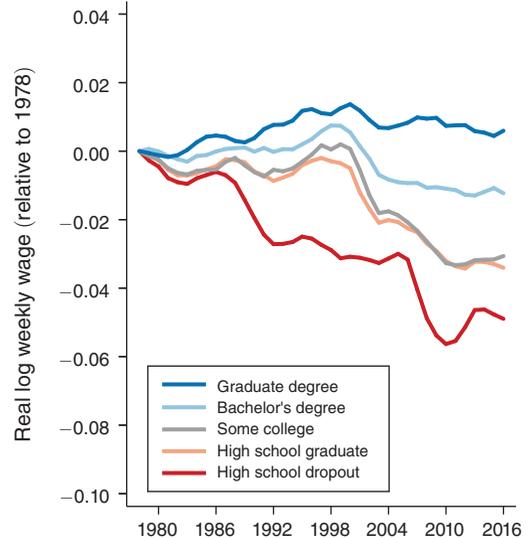


FIGURE 6. REAL LOG WAGE GROWTH BY EDUCATION GROUP, 1978 TO 2016: OBSERVED VERSUS BETWEEN OCCUPATION REALLOCATION COMPONENT, 1970–2016

Notes: Data source is as in Figure 1. Each data point is a three-year centered average of the mean log wage for the relevant education group using earnings years 1977–2017. Series are normalized to zero in 1978. See equation (1) and body text for details of construction.

the two series: the y-range of the observed series is *five* times the range of the counterfactual series, meaning that the counterfactual series is qualitatively on target but quantitatively way off. Hence, changes in occupational composition alone cannot explain (most of) the evolution of between-group wage inequality in these three decades. This conclusion should not come as a surprise. An occupation is not a labor market—hence changes in labor demand in one occupation should affect wages in other occupations. A general decline in demand for non-college workers in production and clerical jobs—occupations that employed 43 percent of non-college workers in 1978—should reduce non-college wages both by inducing displaced workers to take lower paid service and manual jobs and by placing downward supply pressure on the equilibrium wage of non-college workers generally. Thus, this partial equilibrium exercise arguably provides a (loose) lower bound on the extent to which the reallocation of non-college labor from middle- to low-skill occupations could potentially affect wages for non-college workers.

This is not the end of the story, however. The exercise above implicitly assumes that the decline of middle-skill occupations has occurred at the average (log) wage level within each occupation-education-gender group.¹³ While this is a reasonable baseline assumption, it would be violated if the marginal declining (or growing) job within an occupation differed from the average of that occupation. One scenario in which this would arise is if the decline of middle-wage occupations were particularly concentrated in cities and metro areas where wage levels are consistently higher. As I show in the next two sections, this scenario, in large part, is what has played out over the last four decades.

II. The Geography of Polarization

The structure of work differs across places: local labor markets are often specialized in particular industries and services such as manufacturing, education, entertainment, or health-care. One of the key predictors of occupational structure is population density. Some activities

¹³This assumption is built into the ω_{jkt_s} terms in equation (1). Since these terms correspond to the mean (start-of-period) wage in each occupation \times education \times gender cell, they do not admit within-cell wage heterogeneity.

intrinsically take place in low-density areas, such as agriculture. US manufacturing was concentrated in large cities at the start of the twentieth century, and it slowly migrated towards less dense areas as transportation networks improved (Glaeser 2011). Knowledge-intensive industries tend to locate in cities, where educated workers are most prevalent (Glaeser and Mare 2001; Moretti 2004; Berry and Glaeser 2005). These regularities suggest that occupational structure should vary systematically with population density—if so, the expansion or contraction of specific industries and occupations would be expected to have non-neutral impacts on the structure of occupations across urban, metropolitan, suburban, and rural areas. As it turns out, these simple regularities apply with startling clarity to the evolving occupational geography of US labor markets.

Figure 7 presents a bin-scatter depicting the aggregate relationship between population density and occupational structure at the level of commuting zones (CZs) covering the contiguous US states over the course of the five-and-a-half decades between 1970 and 2015. The three panels of this figure report the CZ-level share of employment among working-age adults in the three broad occupational categories referenced previously: services, transportation, laborer, and construction workers (“low-skill”); clerical, administrative support, sales, and production workers (“mid-skill”); and professional, technical, and managerial workers (“high-skill”).¹⁴ The run variable in this and subsequent figures is the natural log of population density (i.e., number of residents divided by CZ land area). To enforce consistency of CZ rankings across time permits, I use each CZ’s population density in 1970 as the run variable throughout.¹⁵ The data are weighted by the count of working-age adults in each CZ. Hence, each plotted point in the bin-scatter represents approximately 5 percent of all workers in each year.

¹⁴In each panel, I subtract off the overall working-age mean share of employment in the relevant occupational category in 1970, so plotted points correspond to the CZ’s share of employment in the occupational cluster relative to the aggregate mean share in that cluster in 1970.

¹⁵Because the correlation between CZ population density exceeds 0.95 across all decades between 1970 and 2015 (and exceeds 0.91 between 1950 and 2015), using the contemporaneous versus fixed density measure has no meaningful effect on these patterns.

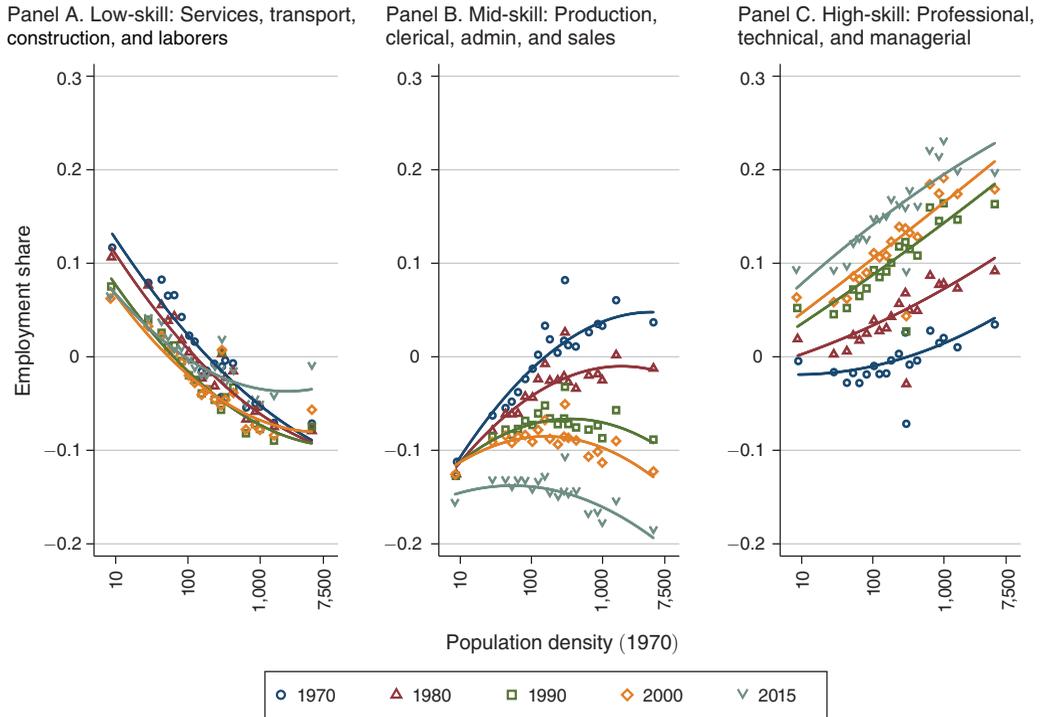


FIGURE 7. OCCUPATIONAL EMPLOYMENT SHARES AMONG WORKING-AGE ADULTS BY COMMUTING ZONE POPULATION DENSITY, 1970–2015: LEVEL RELATIVE TO 1970 MEAN

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of 722 consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

Figure 7 strongly reinforces the conclusion that the decline of middle-skill employment is fully absorbed by a simultaneous rise in high-skill employment—that is, there is essentially no aggregate change in the share of workers employed in traditionally low-skill jobs over the course of 45 years. This figure also conveys three fresh insights. First, while denser CZs have traditionally been more intensive in high-skill work, the level and slope of this density-skill-intensity relationship rose consistently over multiple decades. In 1970, for example, the high-skill occupation share in the densest CZs was about 5 percentage points above the corresponding share in the least dense CZs. By 2015, this gap had increased to approximately 15 percentage points. Second, the fraction of workers engaged in low-skill occupations has historically been considerably smaller in

high-density CZs, and this gradient has changed little over decades. From 1970 through 2000, the low-skill occupation share was consistently 20 percentage points lower in the most versus least dense CZs, after which the fraction of low-skill work in the densest CZs rose by several percentage points over the next 15 years.

Perhaps most strikingly, Figure 7 reveals that denser CZs were exceptional in the 1970s in having far more middle-skill work than suburban and rural CZs. But this exceptional feature attenuated and subsequently reversed sign over the next four decades. In 1970, there was a strong *positive* density gradient in middle-skill employment, with an urban-rural difference of approximately 15 percentage points in the share of workers engaged in clerical, administrative, sales, and production work. This gradient sharply eroded between 1970 and 1990

and eventually reversed sign—from positive to negative—over the subsequent 25 years. While middle-skill work was differentially present in cities and metro areas in the 1970s, it was differentially absent from those locales 45 years later and less prevalent everywhere in absolute terms.

Since Figure 7 shows that the urban expansion of high-skill occupations has almost precisely offset the urban contraction of middle-skill occupations, this figure appears to convey the march of workers up the occupational skill ladder. We know, however, from Figure 5 that this aggregate pattern masks strong compositional shifts within education groups—specifically, an increase in the share of non-college workers employed in historically low-skill occupations. The two panels of Figure 8 show how this reallocation of labor across occupations has unfolded across geographic areas. The upper panel shows that there has been almost no increase in the share of college adults (i.e., those with some college or more) employed in high-skill professional, technical, and managerial occupations.¹⁶ There was, however, a modest decline of about 8 percentage points over 45 years in the share of college workers working in middle-skill occupations and a commensurate rise in the share working in low-skill occupations; moreover, this shift was most pronounced in denser CZs.

Occupational polarization among non-college workers has, however, been far more dramatic, as shown in the lower panel of Figure 8. At no point from (at least) 1970 forward was there any meaningful density gradient in high-skill work among non-college workers. But there was a steep density gradient in middle-skill work. In 1970, non-college workers in the densest CZs were approximately 25 percentage points more likely to work in middle-skill occupations (and 25 percentage points less likely to work in low-skill occupations) than were non-college workers in low-density CZs. This gradient became shallower and its intercept fell over the ensuing 45 years. By 2015, the *low-skill* employment share among non-college workers was several points *higher* in the most versus least dense CZs, while the middle-skill employment share was correspondingly several points lower.

¹⁶Of course, the fraction of college adults has risen substantially (see Figure 2), so no decline in the prevalence of high-skill work among this group implies a substantial increase in the share of workers in high-skill occupations.

Thus, the decline of middle-skill occupations has meant a profound reallocation of non-college workers in urban and metro areas from middle-skill production and office work to low-skill services, transportation, and laborer occupations. This reallocation has been so sweeping that nothing remains of the density gradient in middle-skill work for non-college workers that was strongly evident just four decades earlier.¹⁷

A. Polarization and Immigration

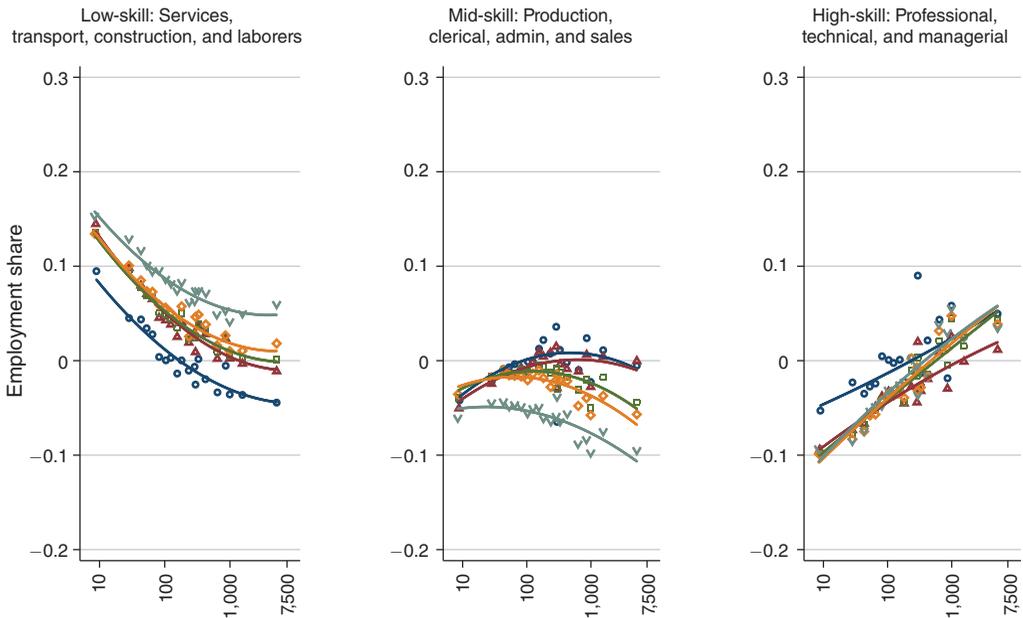
The urban workforce is disproportionately college-educated and foreign-born, and it has become more so over time. Figure 9 illustrates the rising urban gradient in college degree holding that is noted by multiple scholars (Costa and Kahn 2000; Glaeser and Mare 2001; Florida 2002; Moretti 2013; Diamond 2016). In 1970, working-age adults in the densest CZs were approximately 5 percentage points more likely to hold a four-year (or higher) degree than were those in the least dense CZs. This gap rose to 15 percentage points between 1970 and 1990. By 2015, it had risen further to approximately 25 percentage points. There was no analogous urban-rural divergence in the distribution of the *least-educated* adults. While the high school dropout share of the working-age adult population fell by approximately two-thirds between 1970 and 2015 (see Figure 2), the rural-urban gap in the high school dropout share was little changed. In net, the educational distribution in high-density CZs has become increasingly right-skewed but it has not become increasingly left-skewed.¹⁸

Figure 10 depicts the second differentiating feature of urban labor markets alluded to

¹⁷These findings are consistent with Frank et al. (2018), who document that as of 2014, small US cities were substantially less specialized in hard-to-automate managerial and technical professional occupations than are larger cities, and thus face greater potential impacts from automation.

¹⁸Employing the simple college/non-college classification used above (i.e., college refers to some college or above and non-college refers to high school graduate or below), the data suggest even less of an overall educational divergence. The urban-rural divide in working-age college residents rose from approximately 5 percentage points in 1970 to 10 percentage points in 2015, while the urban-rural divide in non-college residents fell by the same amount. This is not a dramatic change.

Panel A. College adults



Panel B. Non-college adults

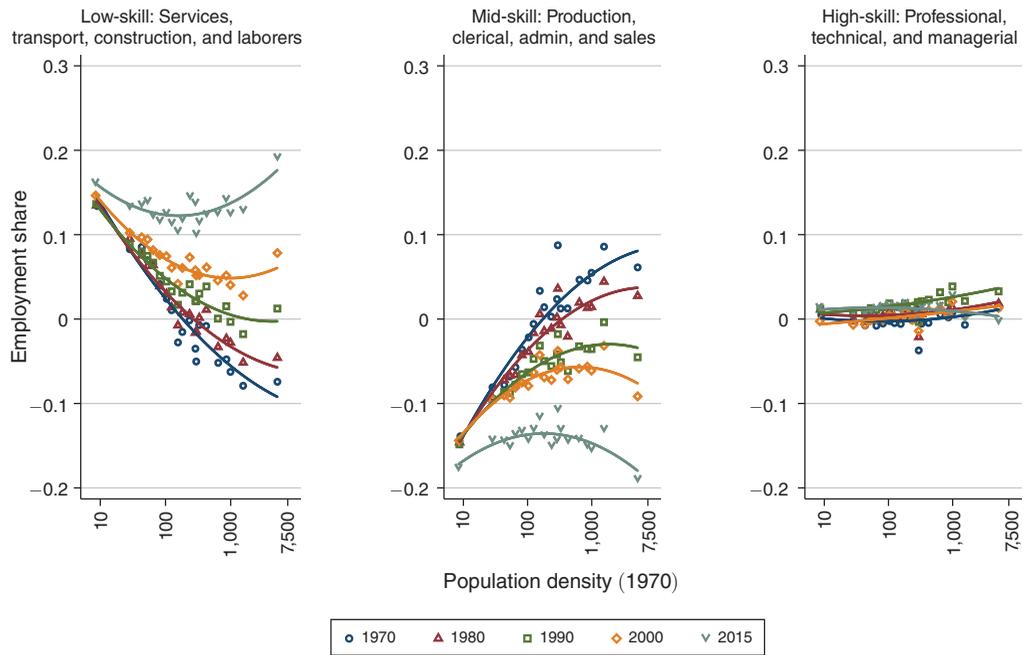


FIGURE 8. OCCUPATIONAL EMPLOYMENT SHARES AMONG (A) COLLEGE ADULTS AND (B) NON-COLLEGE ADULTS BY COMMUTING ZONE POPULATION DENSITY, 1970–2015: LEVEL RELATIVE TO 1970 MEAN

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

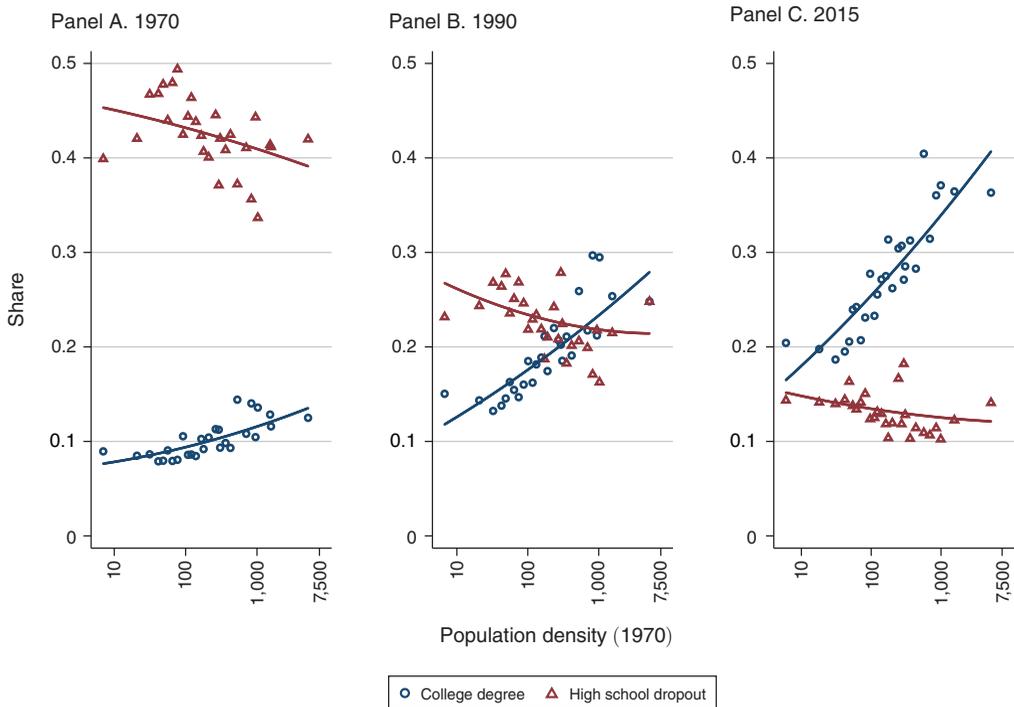


FIGURE 9. COLLEGE DEGREE VERSUS HIGH SCHOOL DROPOUT SHARE OF WORKING-AGE ADULTS, 1970–2015

Note: Figure plots the share of working-age adult residents by CZ who have either four-plus years of college or less than a high school degree.

Source: US Census of Population data for 1970 and 1990 and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Each plotted point represents approximately 3.3 percent of the working-age population in the relevant year.

above: the rise in immigrant intensity. In 2015, the rural-urban gap in the foreign-born share of college adults was approximately 35 percentage points, roughly twice as large as in 1970. Similarly, the rural-urban gap in the foreign-born share of non-college adults was approximately 25 percentage points in 2015, again roughly double the gap in 1970. Foreign-born workers in turn have a bimodal education distribution: they are disproportionately likely to either have completed post-baccalaureate education or to lack a high school diploma.¹⁹

These observations raise the possibility that the shifting density gradient in occupational

structure is in part an artifact of the increasingly bimodal educational and nativity structure of denser CZs. I explore this possibility by plotting changes in occupational composition by population density separately for native-born and foreign-born non-college workers in Figure 11. This figure makes clear that occupational polarization has been equally pronounced among foreign-born and native-born non-college adults. And among both groups, the decline of middle-skill employment and the rise of low-skill employment has been steeper in denser CZs.²⁰ Appendix Figure A1 further shows that the overall patterns of occupational change are highly comparable among

¹⁹Using the pooled 2014–2016 ACS files, I find that 12.2 percent of foreign-born versus 10.7 percent of US born working-age adults had a post-college education in 2015, and 23.1 percent of foreign-born versus 10.5 percent of US born working-age adults lacked a high school diploma.

²⁰Though the fall in the urban gradient in middle skill work has been even steeper for US-born than foreign-born workers, the overall decline is about equally large for both groups.

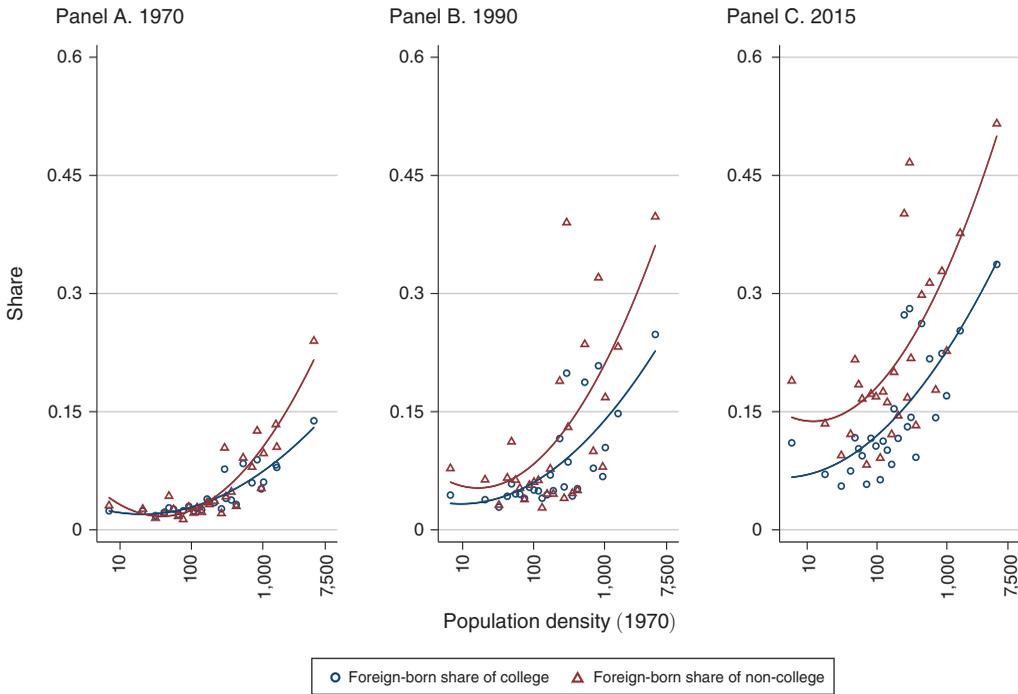


FIGURE 10. FOREIGN-BORN SHARE OF WORKING-AGE ADULTS, 1970–2015

Note: Figure plots the share of working-age adult residents by CZ among college and non-college workers who are foreign born.

Source: US Census of Population data for 1970 and 1990 and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

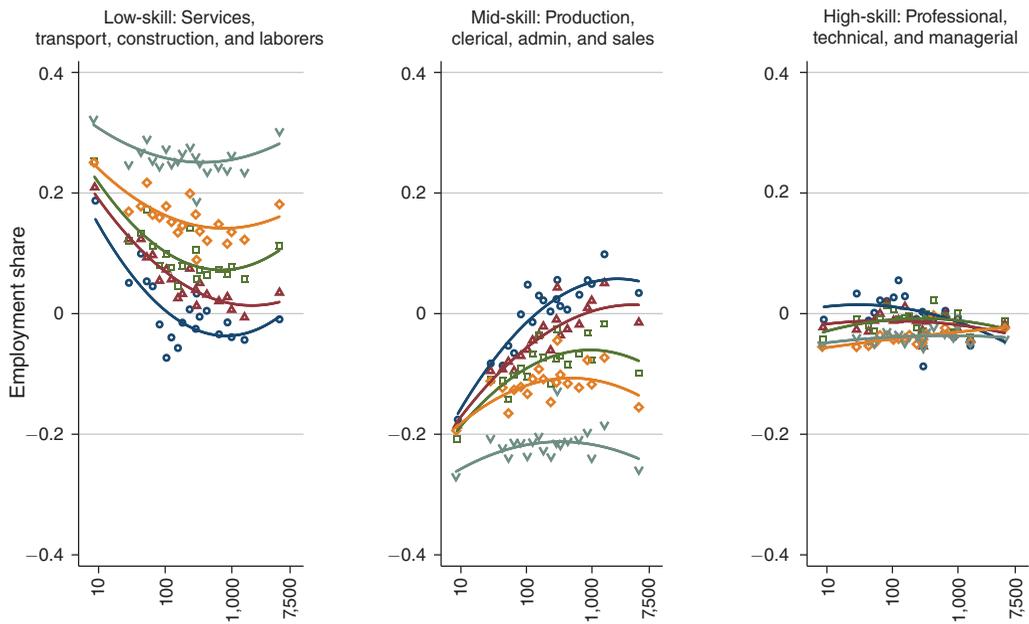
foreign- and native-born college adults. Polarization is therefore not concentrated among foreign-born workers. While one could postulate a more complex dynamic in which rising urban immigration leads to occupational polarization among *both* non-college natives and non-college immigrants, this story runs counter to the influential finding in Peri and Sparber (2009) that rising immigrant penetration catalyzes occupational *upgrading* among similarly educated natives. It therefore seems unlikely that simple supply-side factors such as immigration can explain the polarization of urban employment.

B. The Decline of Urban Production, Clerical, and Administrative Occupations

Much evidence suggests that the pronounced polarization of urban employment originates in substantial part from two secular demand-side

forces: (i) the decline of manufacturing production work in the face of advancing automation and rising trade pressure and (ii) the proliferation of office computing that has hollowed out the ranks of clerical and administrative workers (Autor and Dorn 2013; Autor, Manning, and Smith 2016; Acemoglu and Restrepo 2017). I suspect that the latter force has been particularly significant for non-college workers, who performed the relatively routine subset of office clerical tasks. Both phenomena are visible in Figure 12, which plots the shrinking employment shares of non-college workers in production work and in administrative and clerical work between 1970 and 2015. In 1970, production employment followed an inverted-U in population density—strongly increasing in density through most of the distribution, but tailing off sharply in the densest CZs, reflecting the secular suburbanization of manufacturing (Glaeser

Panel A. Foreign-born non-college adults



Panel B. Native-born non-college adults

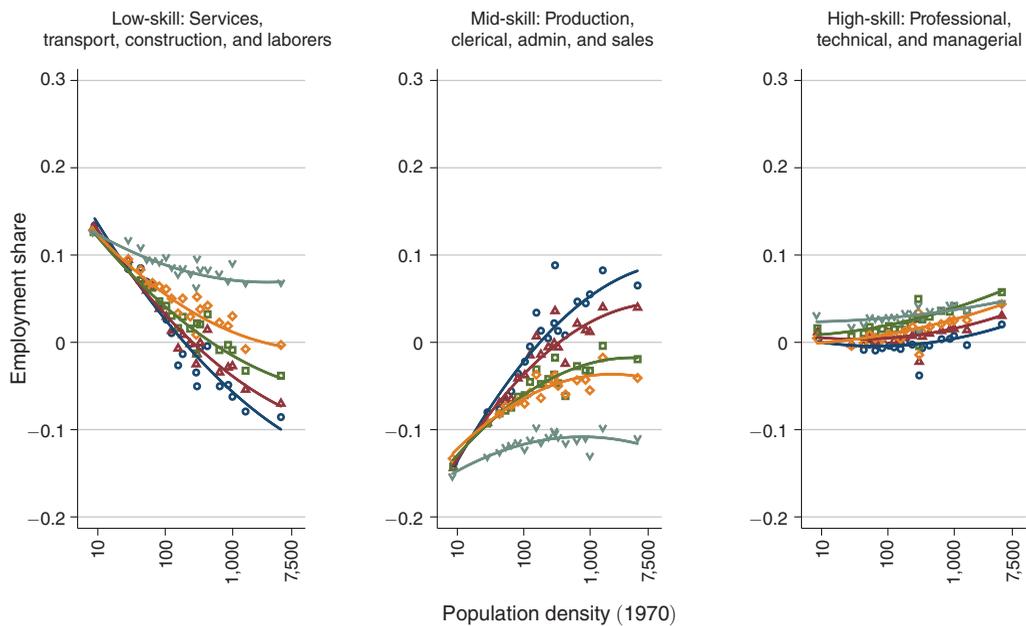


FIGURE 11. OCCUPATIONAL EMPLOYMENT SHARES AMONG (A) FOREIGN-BORN AND (B) NATIVE-BORN NON-COLLEGE ADULTS BY COMMUTING ZONE POPULATION DENSITY, 1970–2015: LEVEL RELATIVE TO 1970 MEAN

Note: Figure plots the share of working-age adult residents by CZ among college and non-college workers who are foreign born.

Source: US Census of Population data for 1970 and 1990 and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

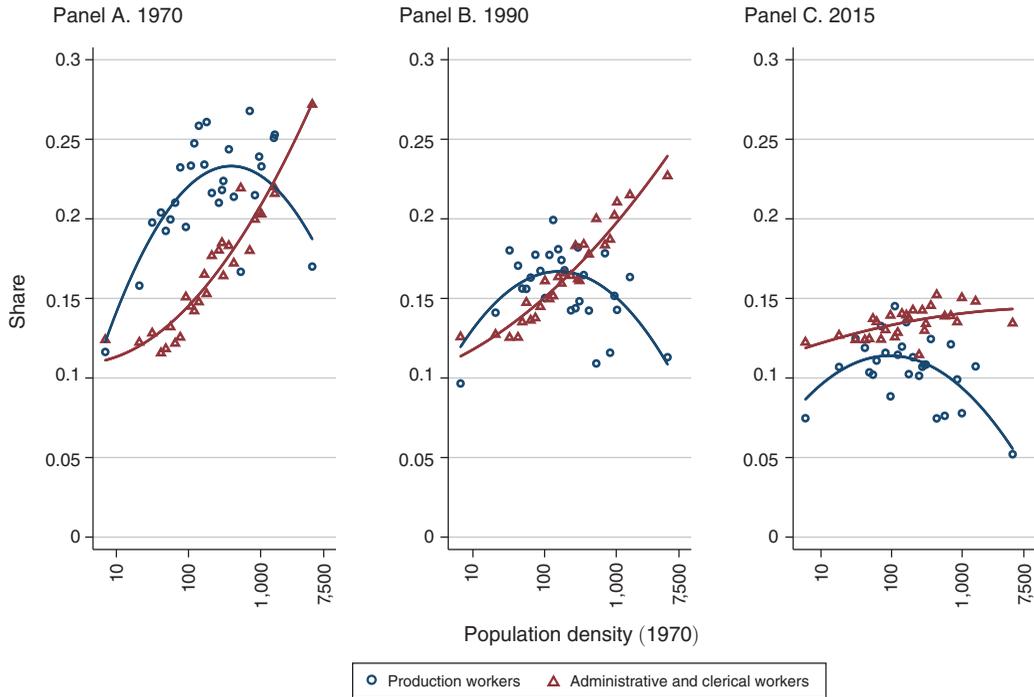


FIGURE 12. PRODUCTION AND ADMINISTRATIVE AND CLERICAL EMPLOYMENT SHARES AMONG NON-COLLEGE ADULTS, 1970–2015

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

2011). Between 1970 and 1990, the production employment share of non-college workers fell by at least 10 percentage points in mid-size metropolitan areas and then fell by approximately that amount again over the next 25 years. By 2015, there was almost no density gradient remaining in production employment: manufacturing was scarce in rural areas (as always), even scarcer in the densest urban areas, and just a few percentage points higher in mid-density CZs.²¹

The analogous pattern of declining non-college clerical and administrative employment begins later in this time interval but culminates even more starkly. In 1970, the share of non-college workers in clerical and administrative occupations was 15-plus percentage

points higher in the most versus least dense CZs. Though not shown in the figure, this relationship was essentially unchanged between 1970 and 1980. With the advent of ubiquitous office computing in the 1980s, clerical and administrative employment among non-college workers fell steeply. This decline was already apparent by 1990 and accelerated thereafter. By 2015, there was almost no remaining (positive) density gradient in office work among non-college workers. Logically, this trend was even steeper among women, with a fall of 25 percentage points in the share of non-college women in office work in the densest CZs between 1970 and 2015 (Appendix Figure A2).

In sum, this evidence underscores how much the secular decline in middle-skill employment has reflected the reversal of a distinctive, long-standing feature of urban versus non-urban areas: the employment of less-educated workers in more skill-intensive occupations. It has, of

²¹ The fall of production employment has been much steeper among non-college men than women, consistent with expectations (see Appendix Figure A2).

course, long been recognized that less-educated workers earn higher wages in urban areas (Glaeser and Mare 2001; Moretti 2004). What was previously unknown is that non-college workers performed distinctly different—more skilled—work in metro and urban labor markets. The erosion of this occupational skill gradient over recent decades, particularly in the 2000s, may in turn augur a fall in the non-college urban wage premium.

III. Polarization and the Urban Wage Premium

It is a robust empirical regularity that urban workers earn more than observably similar non-urban workers (Glaeser and Mare 2001; Moretti 2004; Glaeser and Resseger 2010). Given that land prices are intrinsically higher in dense locations, it is logical that higher urban wages compensate workers for the higher cost of urban living (holding amenities constant). For this to be an equilibrium, however, the productivity of urban workers must also be commensurately higher; otherwise, firms would locate elsewhere. These basic observations imply that some set of agglomerative forces—for example, market thickness, ready exchange of ideas, or external economies of scale—must generate higher productivity in urban areas, hypotheses that have received extensive empirical study (see Glaeser and Gottlieb 2009). While the underlying productive forces behind urban agglomerations are not fully understood, it is easy to imagine how these forces might arise for high-skill workers: In knowledge-intensive work, in-person interactions appear to have few close substitutes, meaning that proximity is critical to productivity (Gaspar and Glaeser 1998; Glaeser and Resseger 2010). It is less obvious, however, why these productivity spillovers accrue to low-skill workers. One possibility suggested by the evidence above is that, in past decades, less-educated workers performed higher-skilled work in urban areas: specifically production, clerical, and administrative jobs. Although the productivity of non-college production and office workers may or may not have been *directly* augmented by their urban siting, their work was necessarily colocated alongside the highly-educated knowledge workers who oversaw it. This geographic complementarity could plausibly have generated a positive occupational and wage density gradient for non-college workers. These observations, if correct, imply that the

secular decline of middle-skill urban occupations could have served to depress the urban wage premium among non-college workers.²² I investigate this possibility here.

A. *The Fading Non-College Urban Wage Premium*

Figure 13 provides strong confirmation that the urban non-college wage premium has in fact declined steeply over the course of several decades, most dramatically after 2000.²³ To my knowledge, this pattern has largely escaped notice in the literature, with the important recent exception of Baum-Snow, Freedman, and Pavan (2018) discussed below.²⁴ While it is tempting to attribute this declining premium to wage structure shifts in urban versus non-urban labor markets, it could alternatively be driven by a number of compositional changes in these markets, including age structure, educational composition of the broad “non-college” category, lingering after-effects of the Great Recession, and immigrant penetration. I explore these confounds in succession, starting with age structure.

Autor and Fournier (2019) report a pronounced inversion of the relationship between population density and population age in the United States over the last six decades. Given that the college/non-college wage premium typically rises over the lifecycle (Card and Lemieux 2001), the shifting age composition of urban versus non-urban CZs could potentially skew raw comparisons of college versus non-college wages. Figure 14 confronts this concern by plotting college and non-college wages for subcategories of prime-age workers ages 25 to 39 and 40 to 54. The falling urban non-college wage

²²One would nevertheless expect a modest urban non-college wage gradient to compensate workers for higher urban living costs, and this appears present in the data.

²³This is opposite in sign to the steeply rising urban wage premium among college workers reported by Diamond (2016) and hence does not reflect a general decline in urban wages.

²⁴Ironically, Baum-Snow, Freedman, and Pavan’s (2018) work had escaped *my* notice prior to obtaining the results above. Given that they seem to run against the grain of recent literature on urban inequality (see for example Diamond 2016, Table 8, indicating that rising college worker productivity in cities raises non-college wages), it is reassuring that this pattern is independently documented elsewhere.

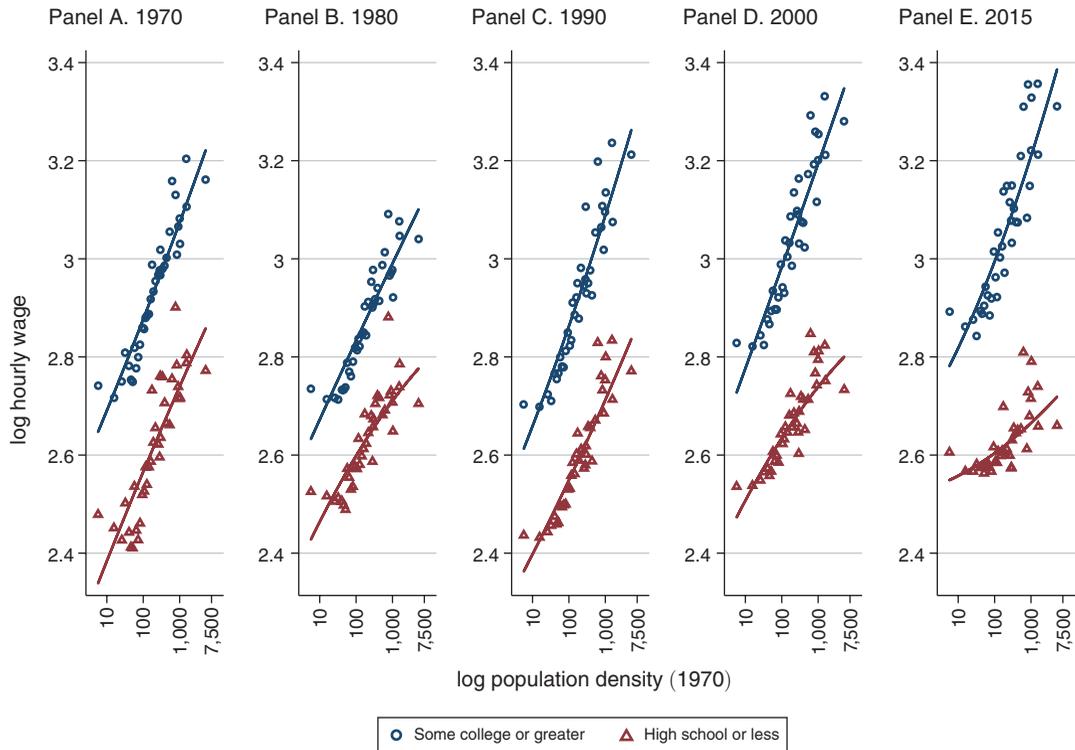


FIGURE 13. REAL LOG HOURLY WAGES OF COLLEGE AND NON-COLLEGE ADULTS, 1970–2015: WORKING-AGE ADULTS

Note: Figure plots real mean log hourly earnings among college and non-college workers in 1970, 1980, 2000, and 2015.

Source: US Census of Population data for 1970, 1980, 1990, and 2000 and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Each plotted point represents approximately 2.5 percent of the working-age population in the relevant year.

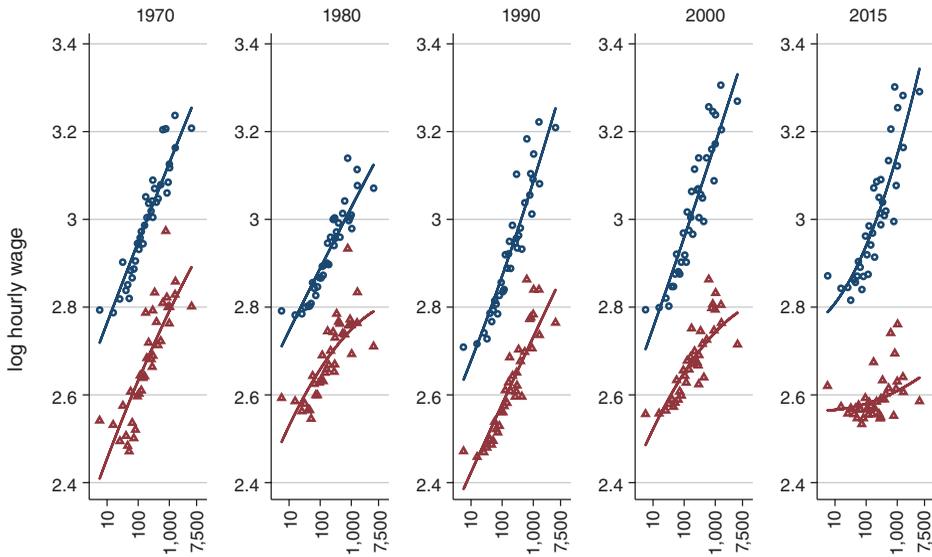
premium is clearly evident within these groups of prime-age workers.

Given the coarseness of the college and non-college education categories, a further concern is that changes in the non-college urban wage gradient could in part reflect changes in the composition of educational attainment within these broad buckets. Figure 15 explores this possibility by plotting wage gradients separately for five detailed education categories: high school dropouts, high school graduates, some college, four-year degree, and post-college educated. The falling urban non-college wage premium is highly visible within subgroups of workers with less than a four-year college education, particularly among high school dropouts and high school graduates, but also among those with some college. Apparently, the declining non-college urban wage premium is not an artifact of the aggregation of educational categories.

A separate concern is that the declining non-college urban wage premium might primarily reflect a lingering consequence of the Great Recession. Figure 15 addresses this concern by presenting panels for both the pre-recession year of 2007 and the post-recession year of 2015. As is evident in the figure and detailed in online Appendix Tables 1 through 3, the fall in the urban wage premium for non-college workers commences well before the Great Recession (indeed, it is visible by 1990) and becomes cumulatively more pronounced.²⁵

²⁵ Wage estimates for the year 1950, reported in the online Appendix tables though excluded from figures for clarity, confirm that the non-college urban wage gradient was *at least* as steep in that decade as the corresponding gradient among college workers. Thus, the fall in the non-college urban wage gradient after 1990 reflects a sharp departure from the prior four decades.

Panel A. Men and women ages 25–39



Panel B. Men and women ages 40–54

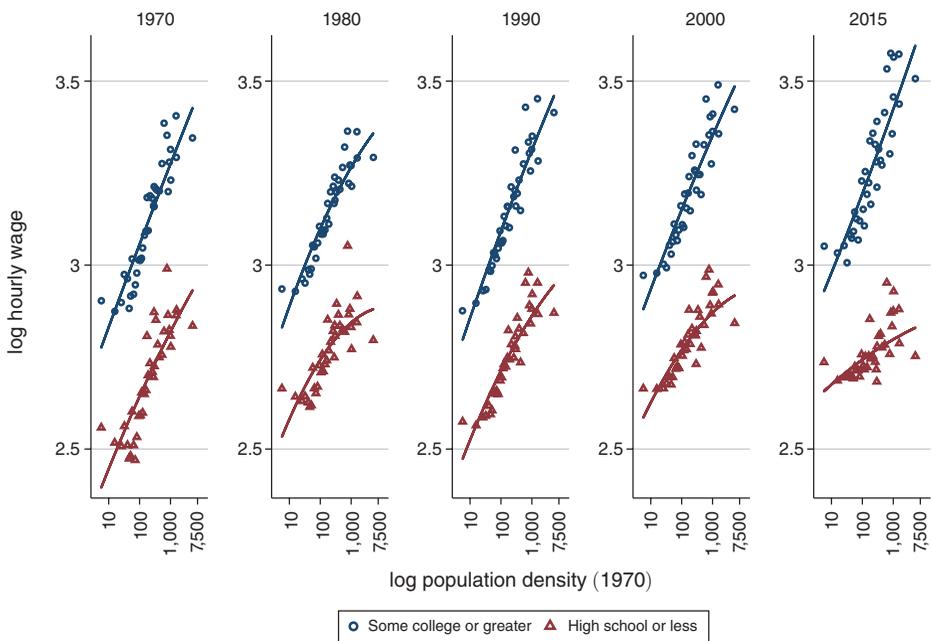


FIGURE 14. REAL LOG HOURLY WAGES OF COLLEGE AND NON-COLLEGE MEN AND WOMEN AGES (A) 25–39 AND (B) 40–54

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 2.5 percent of the working-age population in the relevant year.

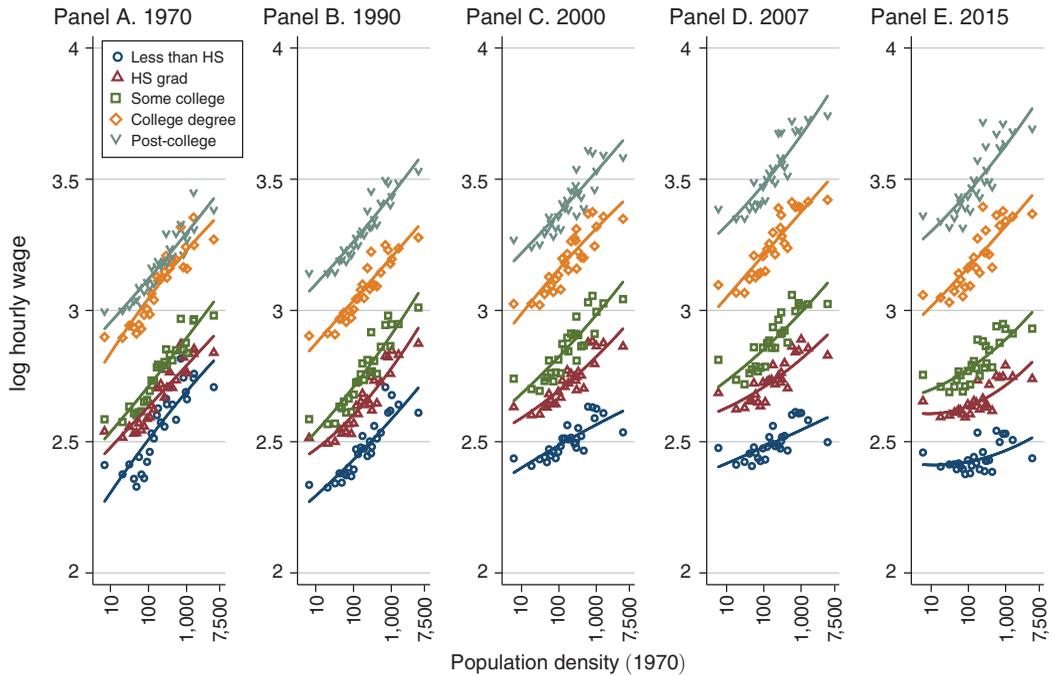


FIGURE 15. REAL LOG HOURLY WAGES BY DETAILED EDUCATION CATEGORY, 1970–2015

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, 2000, and pooled American Community Survey (ACS) data for years 2006 through 2008 and for 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 3.3 percent of the working-age population in the relevant year.

Appendix Figure A3 further documents that the non-college urban wage premium was as at least as steep in the 1950s as in the 1970s, 1980s, and 1990s—so, the post-1990 decline is recent and atypical. Appendix Figures A4 and A5 further document that the decline is equally prevalent among men and women and, in keeping with the evidence above, among both foreign-born and native-born workers (see also online Appendix Tables 4 and 5). I conclude that the falling premium is unlikely to be a compositional artifact.

The declining non-college urban wage premium was (to my knowledge) first reported by Baum-Snow, Freedman, and Pavan (2018) in their study of the causes of rising urban wage inequality from 1980 through 2007.²⁶ Fitting

a structural model that admits both capital skill-complementarity and urban agglomeration, Baum-Snow, Freedman, and Pavan (2018) estimate that cities and skills have become more complementary over time—formally, agglomerative forces for skilled workers have risen—which can rationalize the rising urban college wage premium and falling non-college wage premium. In related work, Giannone (2018) employs a structural model to interpret the decline of regional wage convergence among US cities after 1980. Like Baum-Snow, Freedman, and Pavan (2018), Giannone (2018) concludes that wage divergence among skilled workers across locations is

²⁶Use the same IPUMS source data used here (though for a shorter time interval), Baum-Snow, Freedman, and Pavan (2018) find that not only did the urban wage gradient rise for workers with college or higher education over the last several decades (as is well known), it *fell*

between 1990 and 2000 among workers with high school or lower education. Their measure of economic geography and urban/non-urban differs slightly from those used here. Baum-Snow, Freedman, and Pavan (2018) use Core-Based Statistical Areas rather than Commuting Zones as geographies, and they measure urbanicity using CBSA log population rather than CZ log population *density*.

driven by rising agglomerative forces for skilled labor that, in the model, interact positively with skill-biased technical change.²⁷

The evidence here complements these inferences by identifying a reason why agglomerative forces may have diverged among skill groups—specifically, the withering of the urban occupational skill gradient among non-college workers. This gradient was a vital feature of US economic geography in 1970. It declined modestly in that decade, eroded far more rapidly in the 1980s and 1990s, and was completely absent by 2015.²⁸ Distinct from earlier decades, non-college workers now perform essentially the same work in both urban and non-urban labor markets: custodial work, food services, protective services, recreation, and health services; transportation services; and laborer occupations. The agglomerative forces that plausibly contributed to the steep non-college urban wage gradient in earlier decades appear to have substantially attenuated.²⁹

One supplementary piece of evidence that lends support to this interpretation is the evolution of wages among non-college workers in high-, medium-, and low-skill occupations in the intervening decades. As shown in Appendix Figure A6, the urban wage gap between high- and medium-skill occupations expanded over recent decades while, simultaneously, the urban wage gap between medium- and low-skill occupations contracted, especially for men. These patterns are consistent with falling relative demand for middle-skill work in cities and metro areas. This suggestive evidence should be read cautiously, however. As noted above, an occupation is not a labor market. Similarly, urban and non-urban wage structures are not independent since they

are linked by worker migration and firm arbitrage. It would be premature to conclude that the collapse of middle-skill urban employment explains the flattening of the urban non-college wage premium. Clearly, this is a ripe topic for theoretical and empirical exploration.

B. Accounting for the Geography of Polarization: Wage Implications

I hypothesized above that, by shunting non-college workers into traditionally low-skill occupations, the ongoing encroachment of occupational polarization may in part explain the perplexing fall in real wages of non-college workers over the last three decades. Given the evidence in Section II that occupational polarization has disproportionately occurred in high-wage urban and metro labor markets, it appears likely that the geography of polarization has magnified its wage impacts. I explore this possibility by using the kernel density reweighting technique of DiNardo, Fortin, and Lemieux (1996)—henceforth, DFL—to answer the following question: How would the wages of college and non-college workers have changed between 1970 and 2015 had occupational composition *and* occupational geography evolved as observed while wage levels by occupation and location are held fixed at their 1970 levels.³⁰ This partial equilibrium clearly abstracts from the many economic forces causing wages to evolve within occupations and across places during the intervening decades; concretely, it varies quantities while holding prices fixed. In this application, however, I suspect that this exercise systematically *underestimates* the contribution of occupational change to wage changes by skill group—particularly for the non-college workers who are my focus.

Following DFL, I write the observed wage distribution $f(w)$ in year t_0 as the joint distribution of wages w and covariates x (i.e., education, occupation, sex, and location) integrated over the domain of covariates in year t_0 , denoted as Ω_x :

$$f_{t_0}(w) = \int_{x \in \Omega_x} dF(w, x | t_{w,x} = t_0).$$

³⁰DFL generalizes the venerable Oaxaca-Blinder decomposition from decomposing wage means to decomposing wage distributions.

²⁷Giannone (2018) notes that regional wage divergence has occurred among high but not low-skill workers. She does not report a decline in the urban low-skill wage premium.

²⁸This gradient appears even more pronounced in 1950 than in 1970, but incompatibilities across historical occupational classification schemes reduce my confidence in this inference. Supplementary plots are available from the author.

²⁹These findings may also help to address the puzzling decline of low-skill migration toward high-income US states explored in Ganong and Shoag (2017). Their work attributes this decline to steep increases in housing prices and housing regulations in wealthy regions that increasingly deter new low-skill entrants. The above results suggest a complementary explanation: the flattening urban non-college wage gradient has reduced the incentive for low-skill workers to move to high-income cities.

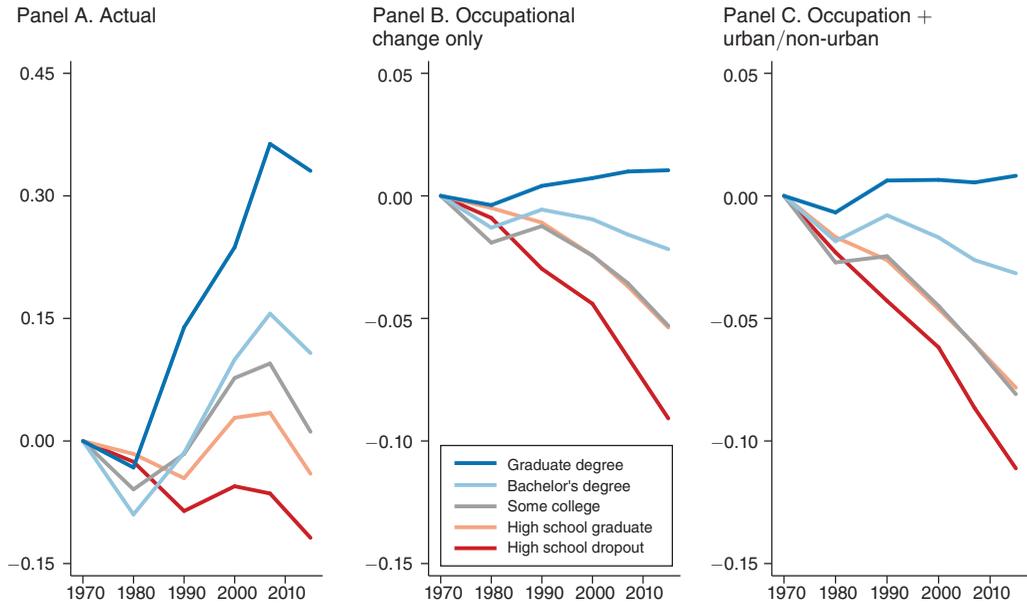


FIGURE 16. OBSERVED AND COUNTERFACTUAL CHANGES IN LOG HOURLY WAGES BY EDUCATION GROUP, 1970–2015

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, 2000, and pooled American Community Survey (ACS) data for years 2006 through 2008 and for 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). The first panel reports cumulative changes in real log hourly wages by education group for years 1970–2015. The second panel reports a DFL reweighting exercise where the 1970 conditional wage distribution is reweighted to match the employment distribution across the 12 occupational categories used above in each subsequent post-1970 period. The third panel repeats this exercise while reweighting the 1970 conditional wage distribution to match the subsequent occupation distribution and the geographic distribution of occupations in each major category (low-, middle-, and high-skill) across more versus less dense commuting zones.

Iterating expectations, this can be rewritten as

$$f_{w_{t_0}}^{x_{t_1}}(w) = \int f(w|x, t_w = t_0) dF(x|t_x = t_0).$$

Here, the distribution of w is conditioned on x and the distribution of x is conditioned on t_0 , as indicated by the subscript and superscript, respectively. Using this identity, I substitute in the x distribution from a subsequent time period, t_1 :

$$\begin{aligned} f_{w_{t_0}}^{x_{t_1}}(w) &= \int f(w|x, t_w = t_0) dF(x|t_x = t_1) \\ &= \int f(w|x, t_w = t_0) \\ &\quad \times \psi_x(x) dF(x|t_x = t_0). \end{aligned}$$

In this expression, the function $\psi_x(x) = dF(x|t_x = t_1)/dF(x|t_x = t_0)$ reweights the distribution of covariates in period t_0 to match those in t_1 (i.e., quantities change). The conditional distribution of wages given covariates x is meanwhile held at its start of period (t_0) level (i.e., prices are held fixed).³¹

Figure 16 reports the results of this reweighting exercise. The first panel reports the familiar pattern of wage inequality in these decades: Real and relative wages of college-educated

³¹In practice, estimating the odds ratio for each value of a high-dimensional vector is fraught with indeterminacies. Following DFL, I invert the problem using Bayes' Rule to write $\hat{\psi}_x = [\Pr(t_x = t_1|x)/\Pr(t_x = t_0|x)] \times [\Pr(t_x = t_0)/\Pr(t_x = t_1)]$, which I estimate with a logit model.

workers fall between 1970 and 1980 (see Katz and Murphy 1992), then rise steeply over the next three and a half decades, until the onset of the Great Recession; real wages of workers with high school or lower education drop steeply during the 1980s and modestly rebound during the 1990s, though high school dropout wages have not regained their 1980 level even 35 years later. The second panel shows the effect of reweighting the 1970 wage distribution to reflect the subsequent occupational distribution during each subsequent decade (1980, 1990, 2000, 2007, and 2015).³² Consistent with the reweighting exercise presented in Section I above, Figure 16 shows that occupational reallocation can proximately account for a substantial share of the fall in non-college wages over this time interval.

The third panel of Figure 16 additionally considers the contribution of occupational geography to changing wage structure. Here, I reweight the 1970 wage distribution to reflect subsequent changes in occupational structure *and* subsequent changes in the locations where high-, medium-, and low-skill work occurs by adding to the reweighting vector interaction terms between 1970 log CZ population density and three broad occupation dummies corresponding to low-, middle-, and high-skill occupations as defined above. Accounting for the geography of polarization substantially adds to the explanatory power of this exercise. As is visible in the figure, it can now account for the entirety of the fall in wages of high school dropouts and high school graduates. Consistent with the fact that polarization has occurred disproportionately among urban, non-college workers, reweighting the 1970 wage distribution to account for the changing geography of occupations magnifies the estimated adverse impact of occupational polarization on wages of non-college workers but has essentially no effect on the wages of college workers.

As with the earlier reweighting exercise reported in Figure 6, this simulation *cannot* account for the rising real wages of highly

educated workers. Mechanically, there has been only modest occupational change among college workers, so occupational reweighting has little impact. Substantively, holding the 1970 wage distribution fixed over the next 45 years omits all productivity growth among highly-educated workers—frequently referred to as “skill biased technical change”—that has occurred in the interim. More fundamentally, this exercise abstracts from all supply and demand forces that affect wage levels within and between occupations. It likely understates the downward pressure that polarization places on wages in low-skill occupations by inducing would-be middle-skill workers to bid for historically low-skill jobs. Similarly, it neglects the impact that rising demand for high-skill workers—seen in the rapid growth of professional, technical, and managerial occupations—exerts on college wages across all occupational categories.

IV. Conclusion: Where Is the Land of Opportunity?

Labor markets in US cities today are vastly more educated and skill-intensive than they were five decades ago. Yet, urban non-college workers perform substantially *less* skilled work than decades earlier, and the once robust non-college urban wage premium has largely flat-lined. The evidence presented here suggests that the hollowing out of middle-skill, non-college blue-collar production and white-collar administrative support jobs, which were once abundant in dense urban labor markets, contributes to these adverse trends in the composition and remuneration of non-college employment.

A critical question that remains unanswered by this evidence is whether a countervailing set of economic forces will soon reverse the decline of middle-skill work, and thereby possibly restore the steep urban occupational skill and wage escalator that afforded greater opportunities to non-college workers in prior decades. Theoretical work by Acemoglu and Restrepo (2018) posits one such set of countervailing forces: increased labor abundance stemming from labor-displacing automation yields an incentive for firms to develop new labor-using tasks that “reinstall” labor’s comparative advantage in a broader range of tasks. Consistent with the arguments of Acemoglu and Restrepo (2018), pioneering work by Lin (2011)

³²I use the same 12 broad occupation categories as applied in Figure 6 and discussed in footnote 12 for each of 5 education groups. This exercise does not distinguish between male and female workers within each occupation by education cell.

documents that “new work”—concretely, the creation of new Census occupational titles—is concentrated in cities (see also Berger and Frey 2016). These findings offer some grounds for optimism about the potential for endogenous technological change to restore labor demand.

Tempering these conclusions, Autor and Salomons (2019) find that new Census occupational titles—as defined by Lin—are themselves strongly polarized among skill categories. Consistent with Lin (2011), Autor and Salomons (2019) identify one rapidly growing set of occupations, which they label “frontier jobs,” that involve producing, installing, maintaining, and deploying new generations of technologies (e.g., robot integration, search engine optimization, radiological medicine). Census data confirm that these occupations are relatively highly paid and are disproportionately held by college-educated men.

Conversely, Autor and Salomons (2019) document a second category of new occupations, also growing in dense urban labor markets, that provide labor-intensive, in-person services to affluent consumers, many of whom reside in high-wage urban labor markets. Examples of these occupations, which Autor and Salomons term “wealth work,” include yoga instruction, sommelier services, pet care, and many forms of personal training and counseling. Distinct from frontier jobs, most wealth work is neither technologically novel nor broadly demanding of technical skills. It is also not highly paid. Workers in wealth work occupations typically earn close to the mean of the wage distribution within their local labor markets and a disproportionate share are women.

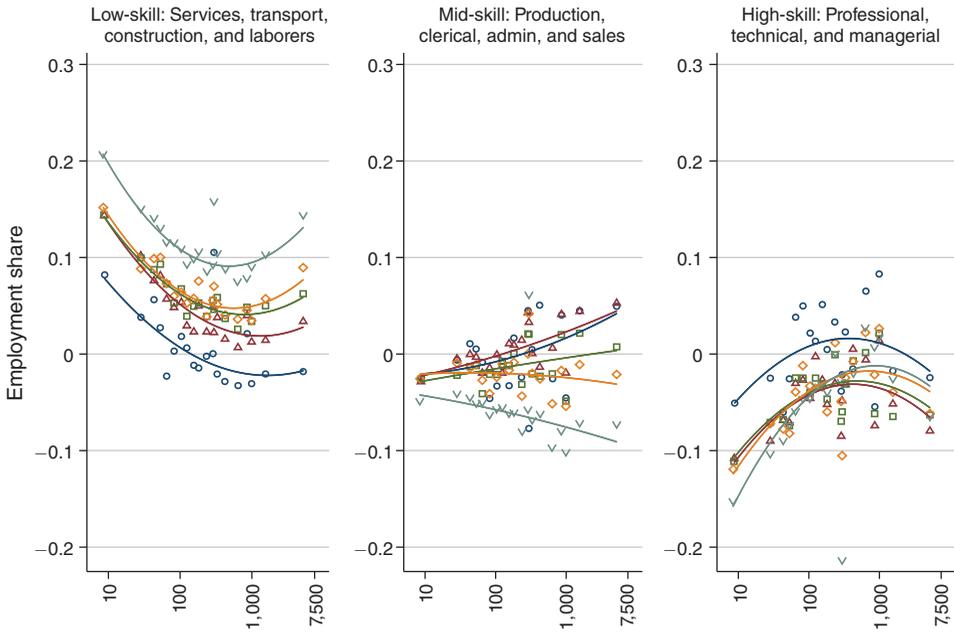
Finally, Autor and Salomons document a third category of new work that involves carrying out nearly-automated tasks that retain a residual set of human components, such as call-center operators, order fulfillment workers, and data entry clerks. These occupations, which Autor and Salomons (2019) dub “last mile” jobs, are less

prevalent in non-urban areas than in urban labor markets, likely because most do not require face-to-face contact with customers. Perhaps unsurprisingly, wage and education levels in last mile jobs are typically considerably below average. In net, these findings support the contention that a disproportionate share of “new work” is generated within dense urban labor markets. But the bifurcated structure of new work does not suggest that a technology-driven “reinstatement” of middle-skill, non-college jobs is currently underway.

Where is the land of opportunity for non-college workers? Recent literature that points to dense urban areas as a growth escalator for US productivity and wages laments, accordingly, the slowing geographic mobility of non-college workers into these locations (see Moretti 2015, Ganong and Shoag 2017, Austin, Glaeser, and Summers 2018). The evidence here points in the opposite direction: the skill escalator that non-college workers once ascended as they entered urban labor markets has lost altitude. And the slowing inflow of non-college workers into urban labor markets may reflect less a failure of arbitrage than a fall in the economic allure that these labor markets once held for less-skilled workers. I view this as a positive development: the slowing migration of non-college workers into high-wage cities should ultimately boost low-skill wages in high-skill labor markets. Simultaneously, the disproportionate aging of the suburban and rural US population during the last four decades (Autor and Fournier 2019) augurs rapidly rising demand for certain labor-intensive, low-skill occupations, such as in-person care, transportation, repair, and other services for the non-urban elderly. This secular demographic change may generate new employment opportunities for non-college workers in low-density locations, and could further reduce or even reverse the long-standing urban non-college wage gradient.

APPENDIX

Panel A. Foreign-born college adults



Panel B. Native-born college adults

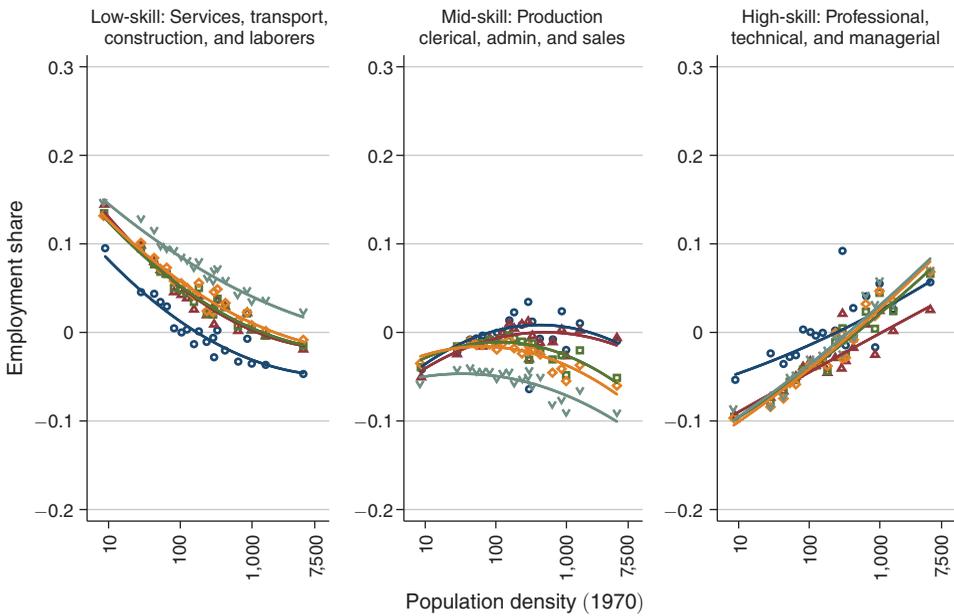


FIGURE A1. LOW-, MIDDLE-, AND HIGH-SKILL OCCUPATION SHARES 1970–2015: RELATIVE TO 1970 MEAN AMONG (A) FOREIGN-BORN AND (B) NATIVE-BORN COLLEGE ADULTS

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

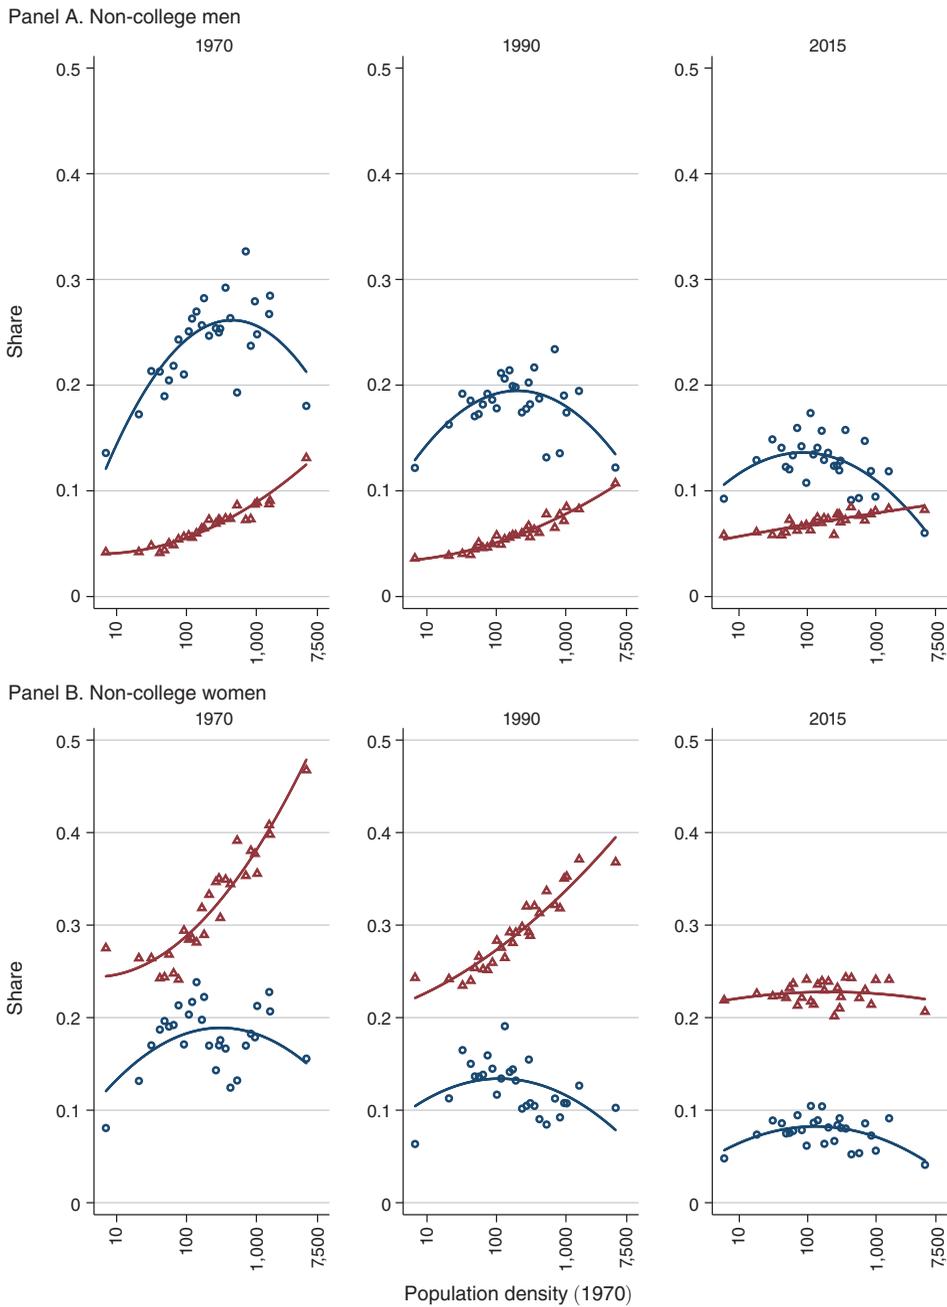


FIGURE A2. PRODUCTION AND ADMINISTRATIVE AND CLERICAL EMPLOYMENT SHARES 1970–2015 AMONG (A) NON-COLLEGE MEN AND (B) NON-COLLEGE WOMEN

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, and 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 5 percent of the working-age population in the relevant year.

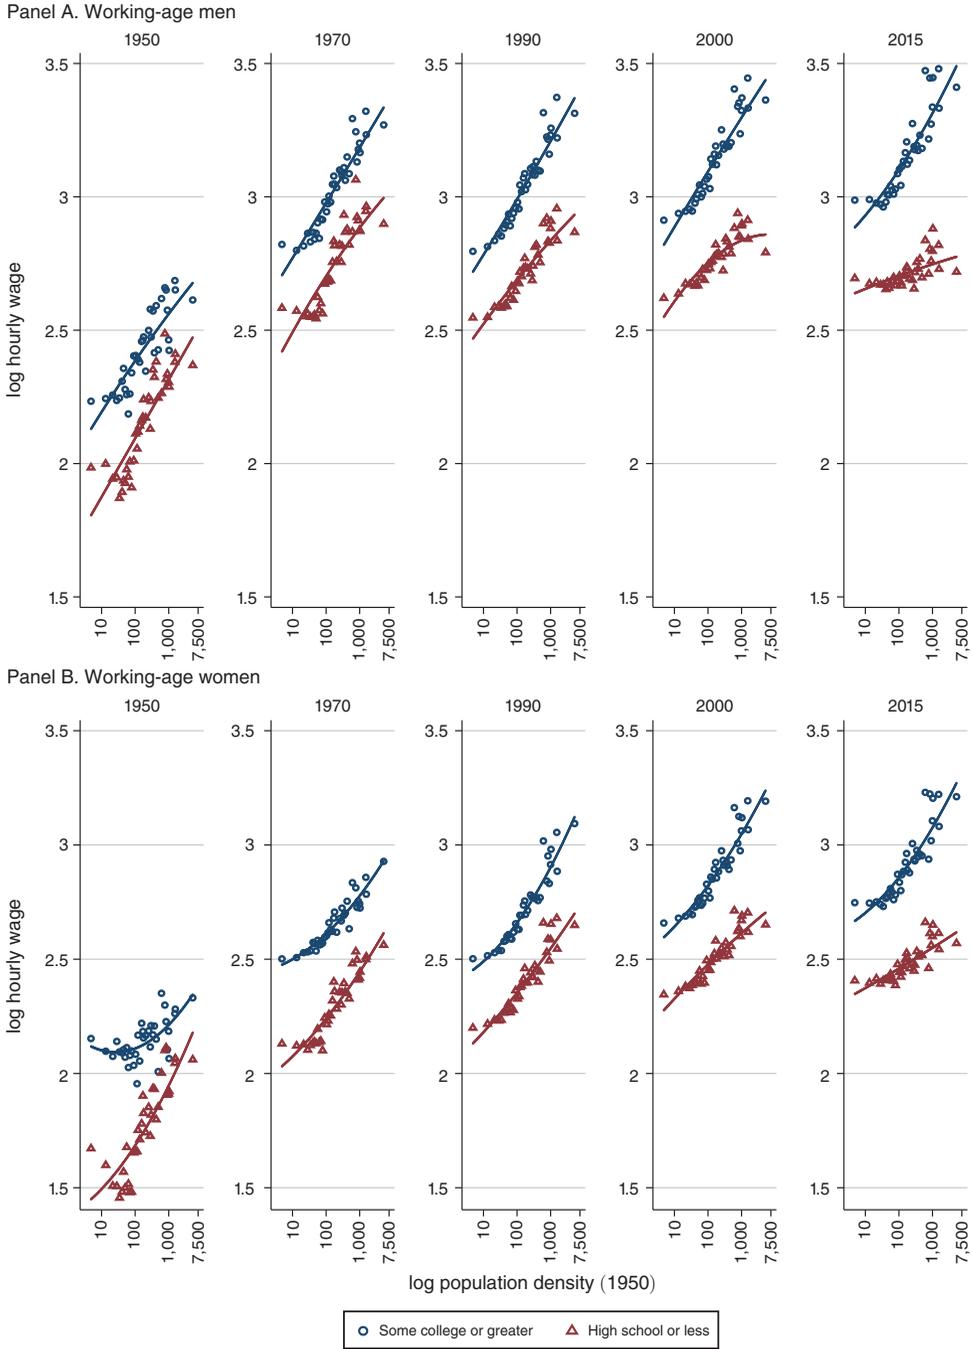
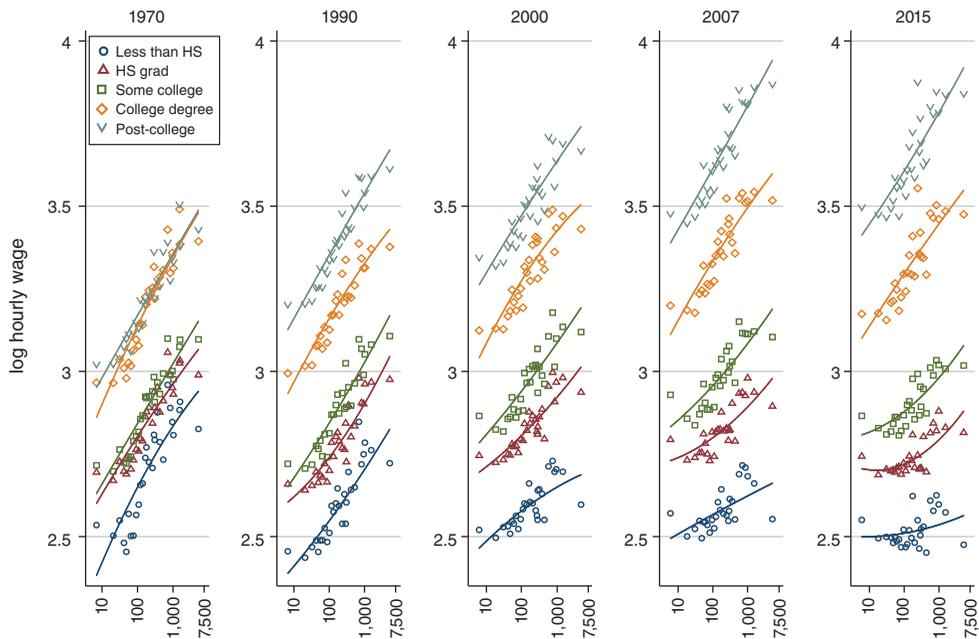


FIGURE A3. REAL LOG HOURLY WAGES OF WORKING-AGE COLLEGE AND NON-COLLEGE WORKERS, 1950–2015: (A) MEN AND (B) WOMEN

Notes: Figure is constructed using US Census of Population data for 1950, 1970, 1990, 2000, and pooled American Community Survey (ACS) data for years 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 2.5 percent of the working-age population in the relevant year.

Panel A. Working-age men



Panel B. Working-age women

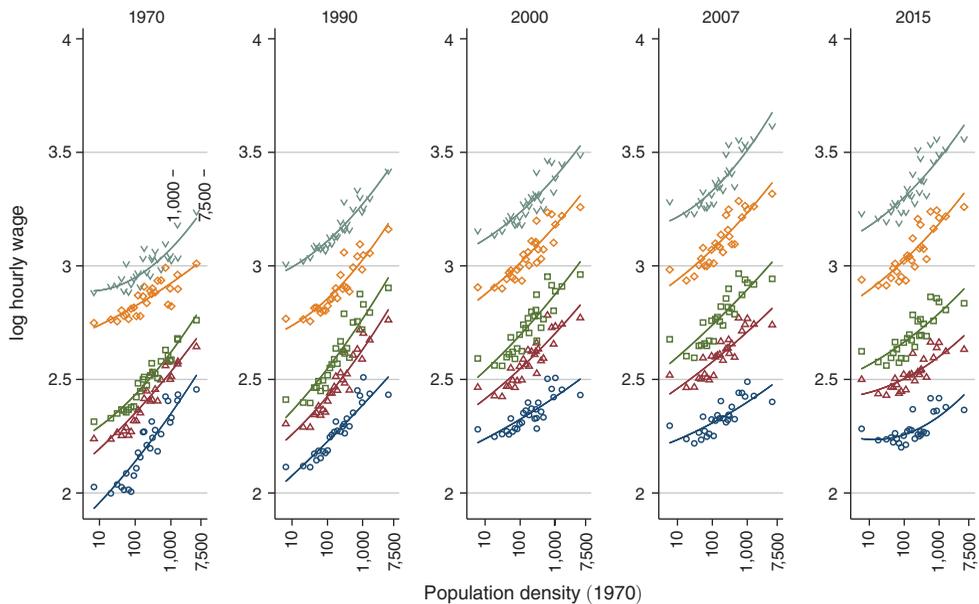
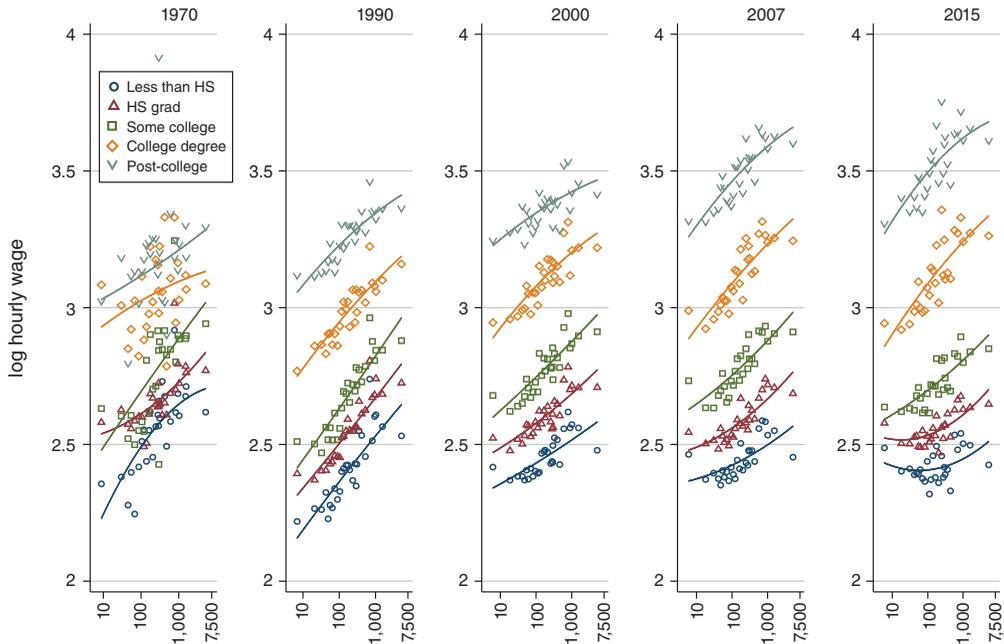


FIGURE A4. REAL LOG HOURLY WAGES BY DETAILED EDUCATION LEVEL, 1970–2015, AMONG WORKING-AGE (A) MEN AND (B) WOMEN

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, 2000, and pooled American Community Survey (ACS) data for years 2006 through 2008 and for 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 3.3 percent of the working-age population in the relevant year.

Panel A. Foreign-born workers



Panel B. Native-born workers

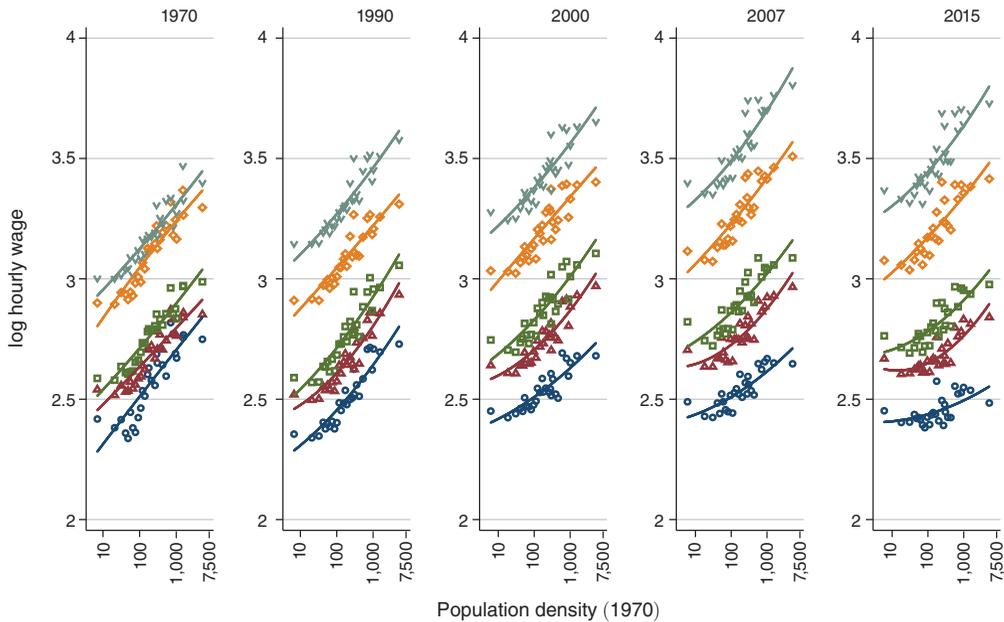
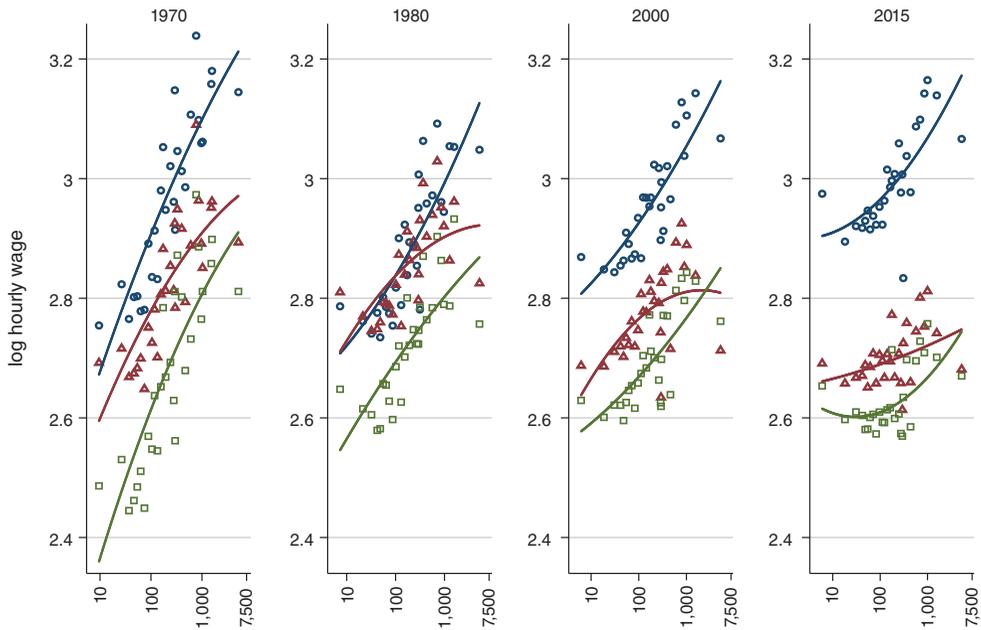


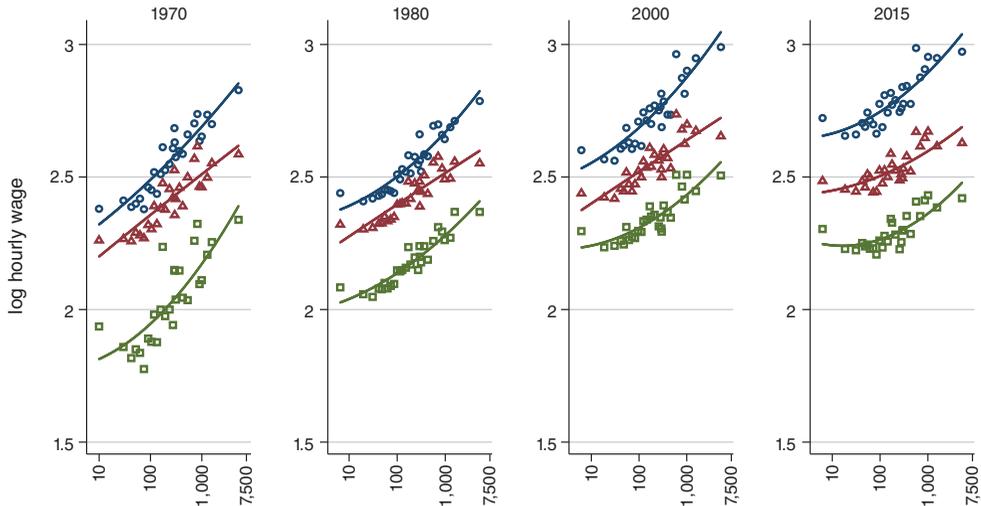
FIGURE A5. REAL LOG HOURLY WAGES BY DETAILED EDUCATION LEVEL, 1970–2015, AMONG WORKING-AGE (A) FOREIGN-BORN AND (B) NATIVE-BORN WORKERS

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, 2000, and pooled American Community Survey (ACS) data for years 2006 through 2008 and for 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 3.3 percent of the working-age population in the relevant year.

Panel A. Non-college men



Panel B. Non-college women



Population density (1970)

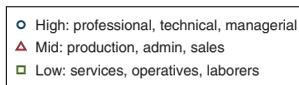


FIGURE A6. REAL LOG HOURLY WAGES BY BROAD OCCUPATION GROUP, 1970–2015, AMONG NON-COLLEGE (A) MEN AND (B) WOMEN

Notes: Figure is constructed using US Census of Population data for 1970, 1980, 1990, 2000, and pooled American Community Survey (ACS) data for years 2006 through 2008 and for 2014 through 2016, sourced from IPUMS (Ruggles et al. 2018). Occupational classifications are harmonized across decades using the classification scheme developed by Dorn (2009) and distilled to the level of consistent local labor markets (AKA, Commuting Zones) following the procedures in Autor and Dorn (2013). Each plotted point represents approximately 3.3 percent of the working-age population in the relevant year.

REFERENCES

- Acemoglu, Daron.** 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1): 7–72.
- Acemoglu, Daron, and David Autor.** 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, Vol. 4B, edited by Orley Ashenfelter and David Card, 1043–1171. Amsterdam: Elsevier.
- Acemoglu, Daron, and Pascual Restrepo.** 2017. "Robots and Jobs: Evidence from US Labor Markets." NBER Working Paper 23285.
- Acemoglu, Daron, and Pascual Restrepo.** 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.
- Aguiar, Mark, and Mark Bilis.** 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105 (9): 2725–56.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad.** 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–1824.
- Alabdulkareem, Ahmad, Morgan R. Frank, Lijun Sun, Bedoor AlShebli, César Hidalgo, and Iyad Rahwan.** 2018. "Unpacking the Polarization of Workplace Skills." *Science Advances* 4 (7): eaao6030.
- Austin, Benjamin, Edward Glaeser, and Lawrence Summers.** 2018. "Jobs for the Heartland: Place-Based Policies in 21st-Century America." *Brookings Papers on Economic Activity* (Spring): 151–240.
- Autor, David H.** 2013. "The 'Task Approach' to Labor Markets: An Overview." *Journal for Labour Market Research* 46 (3): 185–99.
- Autor, David.** 2014. "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent'." *Science* 344 (6186): 843–51.
- Autor, David.** 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3): 3–30.
- Autor, David H., and David Dorn.** 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–97.
- Autor, David H., David Dorn, and Gordon H. Hanson.** 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.
- Autor, David H., David Dorn, and Gordon H. Hanson.** 2016. "The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade." *Annual Review of Economics* 8 (1): 205–40.
- Autor, David, and Juliette Fournier.** 2019. "No Country for Young Men: The Inversion of the Rural-Urban Age Gradient in the United States, 1950–2010." <http://web.mit.edu/dautor/www/Autor-Fournier-AgeInversion.pdf> (accessed March 23, 2019).
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney.** 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96 (2): 189–94.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney.** 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90 (2): 300–323.
- Autor, David H., Alan Manning, and Christopher L. Smith.** 2016. "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment." *American Economic Journal: Applied Economics* 8 (1): 58–99.
- Autor, David, and Anna Salomons.** 2019. "New Frontiers: The Evolving Content and Geography of New Work in the 20th Century." <http://web.mit.edu/dautor/www/Autor-Salomons-NewFrontiers.pdf> (accessed March 23, 2019).
- Baum-Snow, Nathaniel, Matthew Freedman, and Ronni Pavan.** 2018. "Why Has Urban Inequality Increased?" *American Economic Journal: Applied Economics* 10 (4): 1–42.
- Berger, Thor, and Carl Benedikt Frey.** 2016. "Did the Computer Revolution Shift the Fortunes of U.S. Cities? Technology Shocks and the Geography of New Jobs." *Regional Science and Urban Economics* 57: 38–45.
- Berry, Christopher R., and Edward L. Glaeser.** 2005. "The Divergence of Human Capital Levels Across Cities." *Papers in Regional Science* 84 (3): 407–44.
- Blundell, Richard, Robert Joyce, Agnes Norris Keiller, and James P. Ziliak.** 2018. "Income Inequality and the Labour Market in Britain and the US." *Journal of Public Economics* 162: 48–62.
- Card, David.** 1996. "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis."

- Econometrica* 64 (4): 957–79.
- Card, David, and Thomas Lemieux.** 2001. “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis.” *Quarterly Journal of Economics* 116 (2): 705–46.
- Costa, Dora L., and Matthew E. Kahn.** 2000. “Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990.” *Quarterly Journal of Economics* 115 (4): 1287–1315.
- Diamond, Rebecca.** 2016. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000.” *American Economic Review* 106 (3): 479–524.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux.** 1996. “Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach.” *Econometrica* 64 (5): 1001–1044.
- Dorn, David.** 2009. “Essays on Inequality, Spatial Interaction, and the Demand for Skills.” PhD diss., University of St. Gallen.
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg.** 2009. “Revisiting the German Wage Structure.” *Quarterly Journal of Economics* 124 (2): 843–81.
- Florida, Richard L.** 2002. *The Rise of the Creative Class*. New York: Basic Books.
- Frank, Morgan R., Sun Lijun, Cebrian Manuel, Youn Hyejin, and Rahwan Iyad.** 2018. “Small Cities Face Greater Impact from Automation.” *Journal of The Royal Society Interface* 15 (139): 20170946.
- Ganong, Peter, and Daniel Shoag.** 2017. “Why Has Regional Income Convergence in the U.S. Declined?” *Journal of Urban Economics* 102: 76–90.
- Gaspar, Jess, and Edward L. Glaeser.** 1998. “Information Technology and the Future of Cities.” *Journal of Urban Economics* 43 (1): 136–56.
- Giannone, Elisa.** 2018. “Skill-Biased Technical Change and Regional Convergence.” Unpublished.
- Glaeser, Edward.** 2011. *Triumph of the City*. New York: Penguin.
- Glaeser, Edward L., and Joshua D. Gottlieb.** 2009. “The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States.” *Journal of Economic Literature* 47 (4): 983–1028.
- Glaeser, Edward L., and David C. Mare.** 2001. “Cities and Skills.” *Journal of Labor Economics* 19 (2): 316–42.
- Glaeser, Edward L., and Matthew G. Resseger.** 2010. “The Complementarity Between Cities and Skills.” *Journal of Regional Science* 50 (1): 221–44.
- Goldin, Claudia, and Lawrence F. Katz.** 2008. *The Race Between Education and Technology*. Cambridge, MA: Belknap.
- Goos, Maarten, and Alan Manning.** 2007. “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain.” *Review of Economics and Statistics* 89 (1): 118–33.
- Goos, Maarten, Alan Manning, and Anna Salomons.** 2014. “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring.” *American Economic Review* 104 (8): 2509–26.
- Katz, Lawrence F., and Kevin M. Murphy.** 1992. “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *Quarterly Journal of Economics* 107 (1): 35–78.
- Lee, David S.** 1999. “Wage Inequality in The United States During The 1980s: Rising Dispersion Or Falling Minimum Wage?” *Quarterly Journal of Economics* 114 (3): 977–1023.
- Lemieux, Thomas.** 2006. “Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?” *American Economic Review* 96 (3): 461–98.
- Lin, Jeffrey.** 2011. “Technological Adaptation, Cities, and New Work.” *Review of Economics and Statistics* 93 (2): 554–74.
- Mazzolari, Francesca, and Giuseppe Ragusa.** 2013. “Spillovers from High-Skill Consumption to Low-Skill Labor Markets.” *Review of Economics and Statistics* 95 (1): 74–86.
- Michaels, Guy, Ashwini Natraj, John Van Reenen, and John Van Reenen.** 2014. “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 Years.” *Review of Economics and Statistics* 96 (1): 60–77.
- Moretti, Enrico.** 2004. “Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data.” *Journal of Econometrics* 121 (1–2): 175–212.
- Moretti, Enrico.** 2013. “Real Wage Inequality.” *American Economic Journal: Applied Economics* 5 (1): 65–103.
- Moretti, Enrico.** 2015. “Are Cities the New Growth Escalator?” In *The Urban Imperative: Towards Competitive Cities*, edited by Edward L. Glaeser and Joshi-Ghani, 116–148.

- New Delhi: Oxford University Press.
- Peri, Giovanni, and Chad Sparber.** 2009. "Task Specialization, Immigration, and Wages." *American Economic Journal: Applied Economics* 1 (3): 135–69.
- Piketty, Thomas, and Emmanuel Saez.** 2014. "Inequality in the Long Run." *Science* 344 (6186): 838–43.
- Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek.** 2018. IPUMS USA: Version 8.0. type: dataset.
- Tinbergen, Jan.** 1974. "Substitution of Graduate by Other Labour." *Kyklos* 27 (2): 217–26.
- Weil, David.** 2014. *The Fissured Workplace*. Cambridge, MA: Harvard University Press.