

NBER WORKING PAPER SERIES

THE RISE OF HEALTHCARE JOBS

Joshua D. Gottlieb
Neale Mahoney
Kevin Rinz
Victoria Udalova

Working Paper 33583
<http://www.nber.org/papers/w33583>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2025

We are grateful to Scott Blatte and Thuy Minh Le for exceptional research assistance and to Jon Skinner and conference participants at the North American Winter Meeting of the Econometric Society and the Chicago Booth Health Economics Summit for comments. Gottlieb thanks the Becker-Friedman Institute and the Center for Incentive Design for support; This research was made possible in part by a grant from the Peter G. Peterson Foundation. The statements made and the views expressed are solely the responsibility of the authors. This work received support from the Institute for Humane Studies under grant no. IHS017508 (Gottlieb). The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: P-7505738, Disclosure Review Board (DRB) approval number: CBDRB-FY24-0442). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w33583>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Joshua D. Gottlieb, Neale Mahoney, Kevin Rinz, and Victoria Udalova. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Rise of Healthcare Jobs

Joshua D. Gottlieb, Neale Mahoney, Kevin Rinz, and Victoria Udalova

NBER Working Paper No. 33583

March 2025

JEL No. I11, J21, J31, J44

ABSTRACT

Healthcare employment has grown more than twice as fast as the labor force since 1980, overtaking retail trade to become the largest industry by employment in 2009. We document key facts about the rise of healthcare jobs. Earnings for healthcare workers have risen nearly twice as fast as those in other industries, with relatively large increases in the middle and upper-middle parts of the earnings distribution. Healthcare workers have remained predominantly female, with increases in the share of female doctors offsetting increases in the shares of male nurses and aides. Despite a few high-profile examples to the contrary, regions experiencing manufacturing job losses have not systematically reinvented themselves by pivoting from “manufacturing to meds.”

Joshua D. Gottlieb
University of Chicago
Harris School of Public Policy
1307 E. 60th St.
Chicago, IL 60637
and NBER
jgottlieb@uchicago.edu

Kevin Rinz
Washington Center for
Equitable Growth
740 15th St NW
8th Floor
Washington, DC 20005
krinz@equitablegrowth.org

Neale Mahoney
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305
and NBER
nmahoney@stanford.edu

Victoria Udalova
U.S. Census Bureau
4600 Silver Hill Road
Washington, DC 20233
victoria.m.udalova@census.gov

1 Introduction

The United States has experienced rapid growth in healthcare jobs over the past 30 years. As Figure 1 shows, employment in the healthcare industry grew from 9.3 million in 1990 to 18.1 million in 2022. In 2006, healthcare overtook manufacturing in terms of employment, and in 2009 healthcare overtook retail trade to become the largest industry by employment in the U.S.

This paper documents key facts about the rise of healthcare jobs. The motivation for this analysis is simple: Despite the industry’s size and growth, healthcare workers as a whole have not received as much attention as those in other industries or occupations—most notably manufacturing. We use the Census Bureau’s internal versions of the Decennial Census and American Community Survey (ACS) from 1980 to 2022, which provide us with a large sample, more detailed geographic identifiers, and less earnings top-codings than public data.

The first half of the paper documents trends in healthcare employment and earnings through an occupational lens, focusing on familiar clinical occupations such as physicians, nurses, and aides. We show that employment growth has been fairly uniform across most clinical occupations. The exception is a new category known as midlevels, which includes physician assistants and nurse practitioners. This category—which was too small to be consistently measured prior to 2010—has more than doubled since 2010, growing from 227,000 to 505,000 workers. As of 2022, there were more midlevels than primary care physicians, and midlevels provided more than half of primary care services in the U.S. (HRSA, 2023).

This healthcare employment growth was accompanied by strong earnings growth, especially for nurses and midlevels in the middle and upper-middle parts of the clinical occupational distribution. Specifically, earnings grew nearly twice as fast for healthcare workers as for non-healthcare workers from 1980 to 2022; during this window, average healthcare earnings rose from 4% below to over 14% above the average for non-healthcare workers. While the top percentile of the wage distribution has fared better outside of healthcare, health-

care wages have grown faster for the rest of the distribution, and are particularly strong between the middle and the 95th percentiles. Indeed, with strong employment growth and earnings growth that outpaced the rest of the economy outside the very top, it is reasonable to conclude that healthcare has been a modern middle-class “jobs engine.”

The second half of the paper investigates the types of people and places most involved in healthcare employment growth. Our demographic analysis focuses on female and foreign-born workers, two groups whose share of the overall labor force has swelled during our time period. Our geographic analysis focuses on whether growth in healthcare jobs has offset the much-discussed decline in manufacturing employment.

The female share of the healthcare workforce remained nearly constant at 76% between 1980 and 2022, but there was considerable convergence between men and women within occupations. Although physicians remain majority male, the female share increased substantially. Nurses and aides remain overwhelmingly female, but the male share increased. While the foreign-born share of healthcare workers has on average risen in line with that of the broader economy, this differs by occupation: foreign-born shares are well above the economy-wide average for physicians and aides, but below average for nurses and midlevels.

Demographic changes in the overall labor force have been accompanied by major shifts in the geography of employment, with notable declines in the Rust Belt’s employment and population shares during our time period. A policy-focused literature has argued that cities and regions where manufacturing has declined could reinvent themselves by pivoting from “manufacturing to meds.” This literature emphasizes the role of hospitals as anchor institutions and the recession-proof nature of healthcare employment (*e.g.*, Katz, 2006; Bartik and Erickcek, 2008; Kulkarni and Ross, 2013; Harker, Diamond and Reed, 2022; Rolland, 2012).¹ Prominent examples include Pittsburgh, PA, Cleveland, OH, and Rochester, NY. In their analysis of healthcare employment, Kulkarni and Ross (2013) point out that cities with a

¹Aside from a temporary drop during the Covid-19 pandemic, healthcare job growth has indeed been acyclical. As of October 2022, healthcare employment surpassed its level in March 2020 and resumed its pre-pandemic growth rate.

larger historical manufacturing presence now have greater healthcare employment. However, this relationship could either reflect a manufacturing-to-meds pivot, or could arise from city characteristics that jointly influence manufacturing and healthcare, such as size or location.

We systematically test for a manufacturing-to-meds transition by investigating changes over time, examining whether regions that lost more manufacturing jobs disproportionately gained more jobs in healthcare. To avoid reverse causality and (some) omitted variables, we predict declines in manufacturing employment using manufacturing’s share of the prime-age population in each region in 1980. These baseline shares are highly predictive, with a 10 percentage point increase in baseline share corresponding to a 6 percentage point decrease in manufacturing’s share of the prime-age population from 1980–2022 ($R^2 = 0.7$).² However, these places with more baseline manufacturing only experienced modestly higher healthcare employment growth, with each 10 percentage point increase in baseline manufacturing associated with 0.7 percentage point additional growth in healthcare employment as a share of the prime-age population.

As we show with a simple model, a natural benchmark is that industries absorb manufacturing workers in proportion to their sizes. Our estimates imply that healthcare growth counteracted roughly 11% of manufacturing job losses, not much more than would be expected given its 9.8% population share. Healthcare acts as a larger counteracting force for women and college-educated workers; this latter result is consistent with Glaeser’s (2005) argument that high human capital levels enabled Boston to overcome its manufacturing decline. Nevertheless, in the aggregate, our analysis shows that the few high-profile examples of the manufacturing-to-meds pivot are outliers that do not represent a systematic trend.

Our research is related to recent work that examines trends for individual healthcare

²We think about our analysis as descriptive. Our aim is not to recover the causal effect of manufacturing declines on healthcare employment, but instead to document what happened in places that were highly exposed to manufacturing and as a result lost more manufacturing jobs. Our approach addresses concerns about reverse causality (*e.g.* positive shocks to healthcare drawing people out of manufacturing) and certain types of omitted variables (*e.g.* regional productivity shocks that simultaneously effected manufacturing and healthcare employment). It does not address bias from unobserved factors that are correlated with baseline manufacturing.

occupations. Skinner et al. (2019), Staiger, Auerbach and Buerhaus (2012), and Auerbach, Staiger and Buerhaus (2018) have studied the demographics of the physician, nursing, and midlevel workforces, respectively, focusing on age and region. Chan and Chen (2022), Dillender et al. (2022), Alexander and Schnell (2024), Traczynski and Udalova (2018), and others have studied the impacts of growing use of nurse practitioners for patients’ access and quality of care. Gottlieb et al. (2025) document some forces driving physicians’ labor market outcomes, though over a more recent time period than we consider here. Friedrich and Hackmann (2021) and Gottlieb and Zenilman (forthcoming) study the nursing labor market, while Hackmann et al. (2025) study the equilibrium impacts of long-term care employment. Chandra and Heizlsperger (2023) discuss healthcare employment during the Covid-19 pandemic, while Bronsoler, Doyle and Van Reenen (2022) examine how information technology may affect the healthcare workforce.

Relative to this literature, our contribution is to provide a holistic perspective on the healthcare labor market over the decades that the industry has become increasingly central to the U.S. economy. This industry-level perspective allows us to, for example, document strong growth in the middle of the healthcare workers earnings distribution, chronicle convergence in gender ratios across occupations, and compare patterns in healthcare to those in other industries.

Despite our focus on employment, our paper does not take a normative view about the desirability of the observed trends in the quantity or earnings of healthcare jobs. On this topic, Baicker and Chandra (2012) highlight a key tension between the aim of providing affordable access to care and the healthcare system as a jobs program. Chan and Dickstein (2019) and Cooper et al. (2024) explore the extent to which healthcare policy is motivated by political economy considerations around the disbursement of rents and jobs.

The rest of the paper proceeds as follows. Section 2 briefly describes our data, with additional details available in the appendix. Section 3 presents the results on the rise of middle-class healthcare jobs through an occupational lens. Section 4 investigates the types

of people and places that benefited from these trends. Section 5 provides a short discussion of these results and implications for future work.

2 Data

We analyze employment and earnings using microdata from the long-form Decennial Census and American Community Survey (ACS) from 1980 to 2022. These data are repeated cross-sections drawn from nationally representative surveys of the U.S. population. Respondents report earnings, occupation, weeks worked, weekly hours worked, geographic location, educational attainment, and demographic characteristics, among other variables. We use Census data from 1980, 1990, and 2000, and switch to annual ACS data for 2010 and 2022. We start our sample in 1980, immediately after the 1979 peak in manufacturing employment (Harris, 2020), and end with the latest data available in 2022. Our 1980–2022 window covers much of healthcare’s expansion—it grew from 8.9% to 17.3% of GDP during our period (CMS, 2023). We inflation-adjust all monetary variables to 2022 dollars using the CPI-U.

We use the Census Bureau’s internal versions of these data, which have a larger sample, more granular geographic location, and less earnings top-coding than the public microdata. We define earnings as the sum of wage and salary income, business income, and farm income.³ While self-reported earnings in the Census and ACS suffer from some underreporting, especially for the highest earners, they are the best available data for analyzing occupation- and demographic-related labor market trends consistently over this time period.⁴ To the extent any underreporting is consistent over time, it will not affect our analysis of earnings growth. Still, due to these issues, we will use caution when interpreting earnings growth for the very highest earners. See Appendix A for more details.

³This is analogous to the IPUMS harmonized variable “inearn.”

⁴Gottlieb et al. (2025) compare physicians’ self-reported ACS income to comparable values from administrative tax data, though only using data from 2005–2017. They find that self-reports understate total income by about 30%, mostly driven by business income, with greater underreporting for the highest earners.

3 The rise of middle-class occupations

3.1 Clinical occupations

Occupations are a natural lens for examining trends in employment and earnings in health-care. Clinical occupations are commonly known, codified in state and federal regulations, and well-ordered in terms of training and earnings. We analyze trends over our entire 1980–2022 data period and over the more recent 2010–2022 period, where information on midlevels is consistently recorded.

We focus on workers in four clinical occupational groups, which span the range from some of the highest-earning to lowest-earning occupations in the economy.

- Physicians have a medical degree (MD) or doctor of osteopathic medicine (DO) degree. They have generally attended 4 years of college and 4 years of medical school, and have completed 3 or more years of residency and fellowship training, depending on specialty.
- Midlevels (also known as “mid-level practitioners” or “advanced practice providers”) include Nurse Practitioners (NPs), Physician Assistants, Certified Registered Nurse Anesthetists (CRNAs), and Certified Nurse Midwives (CNMs). They have generally studied nursing for 2–4 years and have completed an additional 1–2 years of training. The classifications and training regimes for midlevels have changed rapidly over the past decade. The Nurse Practitioner occupation was too small to be recorded separately in the Census until 2010. To ensure consistent definitions, we restrict our analysis of midlevels to the 2010–2022 sub-period.
- Nurses include Registered Nurses (RNs) and Licensed Practical Nurses (LPNs). In the early part of our sample, they typically had an associates degree. More recently, a bachelor’s degree in nursing (BSN) or even a master’s degree has become common.
- Aides include nursing, psychiatric, and home health aides, medical assistants, pharmacy aides, and similar supporting occupations. Most aides are not required to have

any postsecondary education, but many employers require or strongly encourage some vocational training for certification (BLS, 2024).⁵

In addition to these four categories, we report statistics on two other occupation groups. The first is an “other clinical” occupational group comprising anyone who provides patient care but doesn’t fall into the four categories above (*e.g.*, technicians and therapists). The second is a “non-clinical” occupational group, which includes anyone employed in a healthcare industry who does not report a clinical job, such as managers, billing specialists, janitors, and security staff.

Some workers in clinical healthcare occupations work outside the healthcare industry. For instance, nurses and aides working in assisted living facilities are coded as working in the social services industry, and nurses working at schools are coded as working in the education industry. For completeness, we also report statistics for this “clinical, non-healthcare industry” category. The number of employees in the healthcare industry is the sum of those in the six occupation groups above minus “clinical, non-healthcare” workers.

3.2 Trends in employment and earnings

We begin by examining trends in employment and earnings over 1980–2022 for the healthcare industry and by occupation group, starting with a summary of the 1980 baseline.

Table 1 shows employment and earnings by industry, occupation, and year. In 1980, the healthcare industry employed 7.3 million workers or 6.8% of the labor force. Healthcare workers earned an inflation-adjusted \$44,000 on average, slightly less than the \$46,000 average for non-healthcare workers. In this baseline period, there were 430,000 physicians, 1.7 million nurses, and 1.6 million aides.⁶ Average annual earnings across the clinical occupation groups ranged from \$187,000 for physicians to \$38,000 for nurses and \$23,000 for aides.

⁵For example, federal law requires home health aides to have 75 hours of training and states require up to 180 hours (PHI, 2016). Medical assistants are not required to have any training by law, but the certification required by many employers can take up to a year to earn (Nurse Journal, n.d.).

⁶In 1980, the only occupation in the midlevel occupation group recorded by the Census was Physician Assistants, which had 30,000 employees.

Figure 2(a) shows the changes in total employment by occupation group. The arrows show the change from 1980–2022, and the right column shows the annual growth rate.⁷ For midlevels, we show the change and annual growth over 2010–2022, since we do not have comprehensive data before 2010. Employment has grown roughly twice as fast for healthcare as for non-healthcare workers (2.1% versus 1.1% per year) since 1980. By 2022, the healthcare industry employed 17.8 million workers and accounted for 10.4% of the labor force, a 3.6 percentage point increase since 1980.

Employment growth was broad-based across clinical healthcare occupations, with fairly similar annual growth rates for physicians (2.0%), nurses (2.1%), and slightly higher growth for aides (2.7%). The outlier is midlevels, where employment has exploded in recent years. Over 2010–2022, midlevel employment more than doubled (from 227,000 to 505,000), and the ratio of physicians to midlevels fell by half (from 3.9 to 1.9 times as many physicians as midlevels).⁸

The growth in healthcare employment was accompanied by rapid growth in healthcare earnings. Panel (b) shows an analogous graph for real earnings. These rose nearly twice as fast for healthcare as for non-healthcare workers (1.1% versus 0.7% per year) since 1980. By 2022, the average healthcare worker earned \$71,000 per year, or 14% more than the average non-healthcare worker (compared to 4% less in 1980).

Healthcare earnings grew markedly faster for midlevels and nurses near the middle of the distribution. From 1980–2022, real earnings grew twice as fast for nurses (1.4%) as for aides (0.7% per year) and physicians (0.7% per year). Over 2010–2022, midlevels were the outlier, with real earnings growth of 0.9% per year, compared to declines for nurses and physicians, and modest growth for aides (0.4% per year).⁹

⁷We calculate annual growth rates assuming continuous compounding. The growth rate of a variable Y from its value Y_1 in year T_1 to Y_2 in year T_2 is $\frac{\ln(Y_2) - \ln(Y_1)}{(T_2 - T_1)}$.

⁸Over 2010–2022, there was particularly strong employment growth in clinical occupations outside of healthcare (3.3% per year; see Table 1). Non-clinical healthcare employment grew at 0.7% per year, slower than physicians (0.9%), aides (1.2%), and nurses (1.6% per year).

⁹Between 2010 and 2022, non-clinical healthcare workers experienced 0.6% annual earnings growth, faster than physicians, aides and nurses (See Table 1).

Part of the faster earnings growth could reflect changes in human capital, as nurses and aides have become more educated over time. To investigate this channel, we conduct a DiNardo, Fortin and Lemieux (1996) decomposition, and compare unadjusted earnings growth with earnings growth reweighted to hold the education distribution fixed at its base period level. Appendix Table B.1 shows that earnings grew approximately one-eighth slower for nurses and aides when we adjust for changes in education. In other words, the vast majority of the earnings growth is not explained by changes in education, at least in an accounting sense.¹⁰

The disproportionately fast earnings growth for nurses and midlevels has striking implications for the overall healthcare wage distribution. Figure 2(c) shows real annual hourly wage growth by percentile of the distribution separately for clinical healthcare workers, all healthcare workers, and non-healthcare workers. To construct this plot, we compute real hourly wages at each percentile of the wage distributions in 1980 and 2022. We then take the difference between these values and annualize that difference.¹¹ We restrict this analysis to full-time workers, dropping those who work fewer than 14 weeks or fewer than 20 hours per week on average.¹²

For non-healthcare workers, real wage growth exhibits the sharply increasing pattern documented in Piketty, Saez and Zucman (2018), with average real wage growth below 0.5% until the 76th percentile and then rising from 1.05% at the 95th percentile to 1.55% at the 99th percentile.

For healthcare workers, the wage growth is more equally distributed, although the plot is still upward-sloping across most of the distribution. For healthcare workers overall, real annual wage growth exceeds 0.5% per year for all percentiles above the bottom 7, and peaks at 1.77% at the 95th percentile. For clinical healthcare workers, growth exceeds 1% for all

¹⁰Reweightings does not change the growth rate for physicians, who consistently have a medical degree.

¹¹In other words, the plot shows average annual changes in the inverse CDF of the log wage distribution, or $\frac{\ln(F_{2022}^{-1}(p)) - \ln(F_{1980}^{-1}(p))}{(2022-1980)}$, where p represents a percentile and $F_t(\cdot)$ is the wage distribution in year t .

¹²Following Piketty, Saez and Zucman (2018), we suppress values below the 5th percentile because the low income levels generate noisy growth measures.

percentiles above the bottom 20, and peaks at 2.07% at the 92nd percentile.¹³

These distributions reinforce the lesson from the occupational comparisons. Employment has grown faster for healthcare than for non-healthcare workers, with the fastest growth for midlevels over 2010–2022. Real wages have grown faster for healthcare (clinical and overall) than non-health workers across the entire distribution except the top 1%.¹⁴ Indeed, the gap between healthcare (clinical or all) and non-healthcare is largest between the 45th and 95th percentiles of the wage distributions (Appendix Figure B.2).

4 Implications for people and places

Who has benefited most from the rise of healthcare jobs? Where has this job growth been largest? This section investigates the implications of healthcare’s dramatic employment growth for major demographic groups and geographic areas.

4.1 Demographic patterns

Two prominent trends in the labor market during our time period have been increases in women and immigrants in the labor force. These developments are particularly salient in healthcare; nursing is a classic pink-collar job, and the role of immigrants in the physician workforce is the subject of ongoing policy debates. This section examines trends in healthcare employment through the lens of gender and immigrant status (measured as being foreign born).¹⁵

Figure 3(a) examines trends in the female share of employment by industry and occupation, with the arrows depicting the change between 1980 and 2022. Across the entire

¹³Different income definitions and sample construction choices yield somewhat different results at the very top of the distribution (see Appendix Figure B.3) but the broad patterns are consistent regardless of these choices.

¹⁴We caveat this point with a reminder that our data are least reliable at the very top of the distribution; see footnote 4.

¹⁵The Census and ACS ask about place of birth but do not enable us to construct a more precise measure of immigration status.

economy, the female employment share rose from 42% to 47%. In the healthcare industry, the female share was 76% in 1980 and remained virtually unchanged over the period.

The stable female share for the healthcare industry masks considerable, if incomplete, convergence across clinical occupations. From 1980–2022, the female share of physicians nearly tripled (from 13.3% to 39.8%), although remaining majority male, while the female shares of nurses and aides declined (from 96.1% to 87.7% and 87.2% to 83.7%, respectively), although remaining predominantly female.

Although healthcare’s female share remained constant, because of the industry’s rapid employment growth and high baseline female share, the healthcare industry absorbed a significant portion of the aggregate increase in female employment during this period. Relative to a counterfactual where healthcare employment remains a constant share of the working-age population, and assigning industries their observed changes in female employment, the increase in healthcare from 7.3% to 10.9% of employment accounts for 3.3 million of the 13.1 million women who found employment due to the rise in the female labor force participation.¹⁶

Figure 3(b) examines trends in the share of foreign-born employment by industry and occupation, displayed in the same manner as Panel (a). Across the entire economy, the share of foreign-born workers nearly tripled from 6.7% to 17.3% over this period. The trend in healthcare is similar, growing from 7.7% to 16.3% of the workforce.

As before, the industry-wide trend masks considerable occupation-level variation. Physicians started the period with a 21.2% foreign-born share, or nearly four times the economy-wide average, and grew by another 7.2 percentage points to 28.4% in 2022. In contrast, midlevels’ foreign-born share reached only 12.3% in 2022, nearly one-third less than the economy-wide average and 57% below physicians’. Nurses and aides both started the period with around a 7% foreign-born share, similar to the overall economy, but have since diverged, with the foreign-born share of aides increasing by nearly twice as much as nurses (16.5 versus

¹⁶Appendix A.4 explains this calculation.

8.8 percentage points).

Taken together, this analysis paints a nuanced picture. While the female share of overall healthcare employment has held steady at 76%, there has been significant but incomplete convergence across occupations. While growth in the foreign-born share of healthcare workers has tracked that for the entire economy, foreign-born shares are well above the economy-wide average for physicians and aides and below average for nurses and midlevels.

4.2 Geographic patterns

Another important change during our time period has been the decline in employment and population in the United States' historical manufacturing regions. A policy-focused literature has argued that these Rust Belt cities could reinvent themselves by pivoting from manufacturing to meds. Policymakers and the press have highlighted Pittsburgh, PA, Cleveland, OH, and Rochester, NY as exemplars of this transition.

There is plausible economic logic for this sort of pivot: demand for healthcare services may be particularly resilient to local economic decline. If factory closures reduce overall demand in a region, local demand for healthcare may grow relative to the demand for other goods and services because healthcare spending is supported by government transfers through Medicare and Medicaid. Beyond this natural reallocation, local efforts to encourage healthcare exports (Katz, 2006; Bartik and Erickcek, 2008) could spur a pivot from manufacturing to healthcare.¹⁷ Indeed, Kulkarni and Ross (2013) point out that healthcare jobs are growing rapidly in large metro areas, and observe a correlation between 1980 manufacturing employment and modern healthcare employment. However, this relationship could reflect city characteristics that jointly influence manufacturing and healthcare.

We systematically test for a manufacturing-to-meds reallocation by examining changes over time, investigating whether regions that lost more manufacturing jobs disproportion-

¹⁷There could also be a reverse-Baumol effect where the decline of manufacturing reduces labor costs, enabling other industries to expand. But it is unclear why these effects would be stronger for healthcare than other industries.

ately gained jobs in healthcare. To conduct our analysis, we partition the country into 428 geographic units, defined as the 381 MSAs (2013 vintage) and the 47 non-MSA regions of states. Unlike our prior analysis, which focused on healthcare workers as a share of all workers, we now compute healthcare workers as a share of the prime-age population, defined as persons aged 25–54.¹⁸ This alternative definition is appropriate for this analysis because people can transition from manufacturing to other industries, to unemployment, or leave the labor force entirely. From 1980 to 2022, healthcare grew by 4.4 percentage points—from 5.4% to 9.8% of the prime-age population.¹⁹

An analysis that simply examines the correlation between healthcare job growth and manufacturing job declines could be contaminated by reverse causality and omitted variables. The reverse causality concern is that positive shocks to the healthcare industry could draw workers away from manufacturing. An omitted variables concern is unobserved regional productivity shocks affecting the productivity of both healthcare and manufacturing. To address these issues, we use the baseline manufacturing share to predict manufacturing declines. We then ask whether places with high predicted declines had disproportionately strong healthcare growth.

Econometrically, one can think of baseline manufacturing as an instrument for manufacturing job declines. That said, while our approach is useful in addressing the concerns above, our aim is not to recover the “causal effect” of manufacturing declines on healthcare employment growth. Our approach does not address bias from unobserved factors that are correlated with both healthcare employment growth and the baseline manufacturing share. For instance, if scale economies used to give large cities an advantage in manufacturing, and since 1980 have generated an advantage in healthcare production, our instrument would not

¹⁸This restricts the numerator and denominator to individuals ages 25–54 and includes individuals who are unemployed or not in the labor force in the denominator. Note that the healthcare industry we consider in the numerator is only the healthcare services industry, *i.e.* those involved in caring for patients, and not upstream industries such as research or manufacturing of pharmaceuticals or medical devices.

¹⁹When examined across these geographic units, there is naturally some variation, with increases exhibiting an interquartile range of 3.2 to 5.5 percentage points. That said, 99% of geographic units experienced positive growth. Appendix Figure B.4 shows this distribution.

address this. We thus view our analysis as documenting what happened to healthcare employment in places that were highly exposed to manufacturing—and, as a result, lost more manufacturing jobs—rather than demonstrating a causal impact of this job loss.

Figure 4(a) plots the change in manufacturing share of the prime-age workforce from 1980 to 2022 against the baseline manufacturing share across our geographic units.²⁰ The slope of -0.595 and small intercept means that regions lost on average 60% of their manufacturing workforce between 1980 and 2022. The tight relationship ($R^2 = 0.7$) indicates that the pattern is fairly consistent, with relatively few regions far from this trend.

Figure 4(b) examines whether healthcare employment disproportionately rose in these areas hardest hit by the loss of manufacturing jobs. The slope of 0.066 indicates that healthcare jobs did grow faster in places that lost more manufacturing jobs, although the magnitude of the slope is hard to interpret in isolation.

What is a reasonable benchmark for the increase in healthcare jobs? Appendix C presents a stylized model of an economy with three industries: manufacturing, healthcare, and “other.” We assume the industries face isoelastic demand curves that are identical except for an industry-specific demand shifter and have Cobb-Douglass production functions that are identical except for industry-specific total factor productivity. Labor and capital are perfectly transferable across industries. In this model, shocks that generate declines in manufacturing employment—due to a reduction in either manufacturing demand or manufacturing productivity—cause workers to reallocate to other industries in proportion to those industries’ sizes. Based on this logic, we consider the healthcare employment share as a benchmark for the share of manufacturing job losses that we would expect healthcare to absorb.

Returning to the data, Figure 4(c) plots the change in healthcare employment against the change in manufacturing employment predicted from the 1980 manufacturing share.

²⁰In compliance with Census Bureau disclosure rules, not all regions are included. The scatterplots include 130 regions, representing 77.5% of the U.S. population. Appendix Table B.2 presents coefficients that include all regions. The patterns are identical, although the magnitudes of the relationships are slightly different.

This regression line shows the ratio of 11.1%, as 0.066 healthcare jobs were gained for every 0.595 manufacturing jobs lost. Since healthcare workers comprise 9.8% of the prime age population in 2022, this indicates that the healthcare industry absorbed manufacturing jobs at a rate only slightly higher than the benchmark in which workers flow to other industries in proportion to their size.

In Appendix Table B.3, we examine this relationship between the decline in manufacturing employment and healthcare employment growth separately by demographic group.²¹ Panel (a) reports the relationships discussed above for the entire sample. Panel (b) through (e) show these facts separately for samples defined by the interaction of race (white and non-white) and sex. For this analysis, the statistics are calculated separately at the demographic group level. For example, the 0.228 manufacturing share for white males in 1980 means that 22.8% of white prime-age males worked in manufacturing in 1980.

Panels (b) and (c) show these patterns for white males and females. In 1980, 22.8% of white males and 9.5% of white females worked in manufacturing. For white males, the regression coefficients (column 3) indicate that for each 10% increase in baseline manufacturing employment, white males lost 7.5 percentage points of their manufacturing jobs while gaining only 0.2 percentage points in healthcare employment, or a roughly 1-to-30 ratio of healthcare job gains per manufacturing loss. For white females, the ratio is much larger. For each 10% increase in baseline manufacturing employment, white females lost 3.9 percentage points of manufacturing jobs while gaining 0.9 percentage points in healthcare, a 1-to-4 ratio. Indeed, for white women, healthcare jobs are offsetting nearly twice as many manufacturing jobs as would be expected based on healthcare’s 13% employment share.

For nonwhite men and women, healthcare jobs play less of an offsetting role. In 1980, 21% of nonwhite males and 11% of nonwhite females were employed in manufacturing. For

²¹These appendix tables were computed based on public Census microdata (Ruggles et al., 2024) due to timing constraints making additional disclosures of internal Census results infeasible. The geographic region is Public Use Microdata Area (PUMA), which are only available in a consistent manner from 1980 to 2011. Thus this analysis ends in 2011. Panel (a) of Appendix Table B.3 replicates the aggregate analysis from Figure 4 on this public sample, and results are nearly identical. This gives us confidence in the results based on public-use microdata samples.

each 10% increase in manufacturing employment at baseline, nonwhite men lost 9 percentage points of manufacturing employment while actually losing jobs in healthcare. For each 10% increase in baseline manufacturing employment, nonwhite females lost 5.9 percentage points of manufacturing jobs while gaining 0.5 percentage points in healthcare. The 1-to-11 ratio is less than healthcare’s 13% employment share for nonwhite females and contrasts with the 1-to-4 ratio for white females.

Appendix Table B.4 repeats this exercise on samples split by education. Among those with less than a bachelor’s degree, for each 10% increase in baseline manufacturing employment, there was an additional 6.1 percentage points in manufacturing jobs losses, compared with only 0.5 percentage points in healthcare job gains (a 1-to-11 ratio). For those with a bachelor’s degree or greater, a 10% increase in baseline manufacturing employment was associated with an additional 2.7 percentage points in manufacturing job losses and 0.6 percentage points in healthcare job gains, a ratio of about one-fifth—or twice as large as would be expected based on healthcare’s 11% employment share for this group. Thus, healthcare jobs played a much larger role in offsetting college-educated workers’ manufacturing job losses. This lines up with Glaeser’s (2005) argument that education was critical to reinventing Boston after its manufacturing decline.

In summary, healthcare employment did not play an outsized role in stabilizing the labor markets suffering from manufacturing decline. On average healthcare offsets about 11% of job losses, or only slightly more than would be expected based on healthcare’s 9.8% population share. Digging in, we find that healthcare “punches above its weight” for white women and the college-educated. For white women, healthcare offsets roughly 1 in 4 manufacturing job losses, and for the college-educated, it offsets 1 in 5 manufacturing job losses; both of these ratios are nearly twice what would be expected based on the healthcare employment shares for these groups. For nonwhites and the non-college-educated, on the other hand, healthcare jobs have not been an important stabilizing factor.

5 Discussion

The descriptive analysis in this paper offers three key findings about the rise of healthcare jobs: the relatively strong growth of earnings in the middle and upper-middle parts of the distribution, including for nurses and midlevels; the partial convergence in gender ratios across clinical occupations; and the scant evidence of a systematic manufacturing-to-meds transition, despite high-profile examples.

There are several potential explanations for the strong employment and earnings growth of middle-class healthcare jobs. Entry barriers that limit the supply of doctors may have increased labor demand for midlevels and nurses, who sometimes substitute for doctors, including by providing more than half of primary care (HRSA, 2023). Technological innovation may have differentially benefited workers in the middle of the distribution. For instance, improvements in diagnostic technology may reduce the demand for doctors with sophisticated diagnostic skill in favor of midlevels and nurses who can administer tests and prescribe treatment based on technologically-determined diagnoses. More research is needed to assess the quantitative importance of these and other theories.

The convergence in gender ratios across occupations is consistent with broader evidence on gender dynamics in U.S. labor force. Like many other elite professions (Ferreira and Gyourko, 2014; Wolfers, 2006; Goldin, 2014; Hsieh et al., 2019), there has been a large increase in the female share of physicians. The increasing male share of nurses and aides is consistent with relaxation of gender sorting across occupations.

Why do the high-profile examples of manufacturing-to-meds transitions not reflect a systemic trend? One possibility is that efforts to pivot were largely superficial. Despite policy white papers and political speeches, limited resources were invested in building up the healthcare industry in declining manufacturing regions. Another potential explanation is that any resources that were dedicated to these efforts had lackluster effects. This could reflect the difficulty of breaking into a new industry, especially one where scale effects and reputation are so important (Dingel et al., 2023).

Stepping back, the rise of healthcare jobs is one of the largest trends in the U.S. labor market in a generation. This shift has been accompanied by patterns that sometimes contrast with, and sometimes reinforce, conventional views. Given the size of the healthcare industry, analysis of the U.S. labor force requires understanding the trends in this industry. This paper provides an initial descriptive analysis of these patterns and raises questions for future work. Understanding why healthcare has bucked broader trends toward wage polarization, and whether it will maintain this pattern of broad-based growth, will be crucial as the industry continues to expand and technology changes.


References

- Alexander, Diane, and Molly Schnell.** 2024. “The Impacts of Physician Payments on Patient Access, Use, and Health.” *American Economic Journal: Applied Economics*, 16(3): 142–177.
- Auerbach, David I, Douglas O Staiger, and Peter I Buerhaus.** 2018. “Growing Ranks of Advanced Practice Clinicians—Implications for the Physician Workforce.” *New England Journal of Medicine*, 378(25): 2358–2360.
- Baicker, Katherine, and Amitabh Chandra.** 2012. “The Health Care Jobs Fallacy.” *New England Journal of Medicine*, 366(26): 2433–2435.
- Bartik, Timothy, and George Erickcek.** 2008. “The Local Economic Impact of “Eds & Meds”: How Policies to Expand Universities and Hospitals Affect Metropolitan Economies.” Brookings. Brookings Institution. Available online at <https://www.brookings.edu/articles/the-local-economic-impact-of-eds-meds-how-policies-to-expand-universities-and-hospitals-affect-metropolitan-economies/>.
- BLS.** 2024. “Home Health and Personal Care Aides.” Bureau of Labor Statistics. Available online at <https://www.bls.gov/ooh/healthcare/home-health-aides-and-personal-care-aides.htm>.
- Bronsoler, Ari, Joseph Doyle, and John Van Reenen.** 2022. “The Impact of Health Information and Communication Technology on Clinical Quality, Productivity, and Workers.” *Annual Review of Economics*, 14: 23–46.
- Chan, David C, and Michael J Dickstein.** 2019. “Industry Input in Policy Making: Evidence from Medicare.” *The Quarterly Journal of Economics*, 134(3): 1299–1342.
- Chan, David C., Jr, and Yiqun Chen.** 2022. “The Productivity of Professions: Evidence from the Emergency Department.” National Bureau of Economic Research Working Paper No. 30608, National Bureau of Economic Research.
- Chandra, Amitabh, and Louis-Jonas Heizlsperger.** 2023. “The great resignation, employment, and wages in health care.” *NEJM Catalyst Innovations in Care Delivery*, 4(6).
- CMS.** 2023. “National Health Expenditure Summary Including Share of GDP, CY 1960–2022.” Centers for Medicare and Medicaid Services. Available online at <https://www.cms.gov/files/zip/nhe-summary-including-share-gdp-cy-1960-2022.zip> (accessed October 4, 2024).
- Cooper, Zack, Amanda Kowalski, Eleanor Neff Powell, and Jennifer D Wu.** 2024. “Politics and health care spending in the United States: A case study from the passage of the 2003 Medicare Modernization Act.” *Journal of Health Economics*, 95: 102878.

- Dillender, Marcus, Anthony T. Lo Sasso, Brian J. Phelan, and Michael R. Richards.** 2022. “Occupational Licensing and the Healthcare Labor Market.” National Bureau of Economic Research Working Paper No. 29665, National Bureau of Economic Research.
- 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.
- Dingel, Jonathan I., Joshua D. Gottlieb, Maya Lozinski, and Pauline Mourrot.** 2023. “Market Size and Trade in Medical Services.” National Bureau of Economic Research Working Paper 31030.
- Ferreira, Fernando, and Joseph Gyourko.** 2014. “Does gender matter for political leadership? The case of US mayors.” *Journal of Public Economics*, 112: 24–39.
- Friedrich, Benjamin U, and Martin B Hackmann.** 2021. “The Returns to Nursing: Evidence from a Parental-Leave Program.” *The Review of Economic Studies*, 88(5): 2308–2343.
- Glaeser, Edward L.** 2005. “Reinventing Boston: 1630–2003.” *Journal of Economic Geography*, 5(2): 119–153.
- Goldin, Claudia.** 2014. “A Grand Gender Convergence: Its Last Chapter.” *American Economic Review*, 104(4): 1091–1119.
- Gottlieb, Joshua D., and Avi Zenilman.** forthcoming. “When Nurses Travel: Labor Supply Responses to Peak Demand for Nurses.” *The Review of Economics and Statistics*.
- Gottlieb, Joshua D., David Hémous, Jeffrey Hicks, and Morten G. Olsen.** 2023. “The Spillover Effects of Top Income Inequality.”
- Gottlieb, Joshua D., Maria Polyakova, Kevin Rinz, Hugh Shiple, and Victoria Udalova.** 2025. “The Earnings and Labor Supply of U.S. Physicians.” *The Quarterly Journal of Economics*.
- Hackmann, Martin B, Jörg Heining, Roman Klimke, Maria Polyakova, and Holger Seibert.** 2025. “Health Insurance as Economic Stimulus? Evidence from Long-Term Care Jobs.” National Bureau of Economic Research Working Paper No. 33429.
- Harker, Patrick, Deborah Diamond, and Davin Reed.** 2022. “Anchor Impact: Understanding the Role of Higher Education and Hospitals in Regional Economies.” Federal Reserve Bank of Philadelphia.
- Harris, Katelynn.** 2020. “Forty years of falling manufacturing employment.” *Bureau of Labor Statistics*, 9(16).

- HRSA.** 2023. “State of the Primary Care Workforce 2023.” Health Resources and Services Administration. Available online at <https://bhw.hrsa.gov/sites/default/files/bureau-health-workforce/data-research/state-of-primary-care-workforce-2023.pdf> (accessed December 7, 2024).
- Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow.** 2019. “The allocation of talent and U.S. economic growth.” *Econometrica*, 87(5): 1439–1474.
- Katz, Bruce.** 2006. “Six Ways Cities Can Reach Their Economic Potential.” Brookings Institution. Available online at <https://www.brookings.edu/articles/six-ways-cities-can-reach-their-economic-potential/>.
- Kulkarni, Siddharth, and Martha Ross.** 2013. “Healthcare Metro Monitor Supplement.” Brookings Institution. Available online at <https://www.brookings.edu/articles/healthcare-metro-monitor-supplement/>.
- McKenna, Laura, and Matthew Haubach.** 2019. “Legacy Techniques and Current Research in Disclosure Avoidance at the U.S. Census Bureau.” U.S. Census Bureau.
- Nurse Journal.** n.d.. “How To Become A Medical Assistant | NurseJournal.org.”
- PHI.** 2016. “Home Health Aide Training Requirements by State.” Available online at <https://www.phinational.org/advocacy/home-health-aide-training-requirements-state-2016/>.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman.** 2018. “Distributional National Accounts: Methods and Estimates for the United States*.” *The Quarterly Journal of Economics*, 133(2): 553–609.
- Rolland, Keith.** 2012. “What Makes Cities Resilient?” Federal Reserve Bank of Philadelphia.
- Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Renae Rodgers, and Megan Schouweiler.** 2024. “IPUMS USA: Version 15.0 Census and ACS.”
- Skinner, Lucy, Douglas O Staiger, David I Auerbach, and Peter I Buerhaus.** 2019. “Implications of an aging rural physician workforce.” *New England Journal of Medicine*, 381(4): 299–301.
- Staiger, Douglas O, David I Auerbach, and Peter I Buerhaus.** 2012. “Registered nurse labor supply and the recession—are we in a bubble.” *New England Journal of Medicine*, 366(16): 1463–1465.
- Traczynski, Jeffrey, and Victoria Udalova.** 2018. “Nurse practitioner independence, health care utilization, and health outcomes.” *Journal of Health Economics*, 58: 90–109.
- Wolfers, Justin.** 2006. “Diagnosing discrimination: Stock returns and CEO gender.” *Journal of the European Economic Association*, 4(2-3): 531–541.

Figure 1: Employment for Selected Industries



Notes: This figure shows employment by year for the five largest industries and the manufacturing industry in the United States. Data come from the Bureau of Labor Statistics's (BLS) Current Employment Statistics (CES). Each line corresponds to a two-digit NAICS sector, as defined by the BLS, except for "health care", which is defined as the "health care and social assistance" sector less "social assistance."


Table 1: Employment and Earnings (1980–2022)

Panel A: Employment (Thousands)						
	(1)	(2)	(3)	(4)	(5)	(6)
	1980	2000	2010	2022	% Change 2010–2022	1980–2022
<i>Industries:</i>						
Labor Force	106,400	140,830	157,700	171,110	8.5	60.8
Total Employment	99,570	133,500	140,870	163,850	16.3	64.6
Non-healthcare	92,300	121,730	125,650	146,070	16.3	58.3
Healthcare	7,271	11,764	15,218	17,780	16.8	144.5
<i>Occupations:</i>						
Physicians	433	728	882	981	11.3	126.9
Midlevels	.	.	227	505	122.3	.
Nurses	1,696	2,902	3,423	4,128	20.6	143.4
Aides	1,580	2,746	4,184	4,836	15.6	206.1
Other clinical	1,297	2,332	2,696	4,029	49.4	210.5
Non-clinical	2,857	4,757	5,979	6,511	8.9	127.9
Clinical, non-healthcare	622	1,760	2,173	3,210	47.7	416.4
Panel B: Real Earnings						
<i>Industries:</i>						
Total Employment	46,058	58,295	59,332	62,934	6.1	36.6
Non-healthcare	46,200	57,720	58,200	61,970	6.5	34.1
Healthcare	44,250	64,250	68,680	70,850	3.2	60.1
<i>Occupations:</i>						
Physicians	186,800	238,800	255,400	254,200	-0.5	36.1
Midlevels	.	.	108,400	120,900	11.5	.
Nurses	38,474	62,107	71,350	69,925	-2.0	81.7
Aides	22,790	30,113	28,962	30,393	4.9	33.4
Other clinical	54,447	74,599	79,977	73,215	-8.5	34.5
Non-clinical	33,887	48,669	54,446	58,655	7.7	73.1
Clinical, non-healthcare	46,709	51,545	51,185	50,875	-0.6	8.9


Notes: This table shows employment and earnings by industry and by occupation. Data come from the 1980 and 2000 Decennial Censuses and the 2010 and 2022 American Community Survey. Healthcare and non-healthcare are defined based on 1990 Census industry codes. Occupations are defined based on harmonized occupation codes. Some clinical occupations include workers in non-health industries. Non-clinical occupations includes all employees in a healthcare industry but not in a clinical occupation. We do not report earnings and employment for midlevels before 2010; nurse practitioners, CRNAs, and nurse midwives were classified as nurses until 2010 changes in Census occupation codes, so the midlevel category is incomplete from 1980–2009. Earnings are inflation-adjusted to 2022 dollars using the CPI-U. DRB number: CBDRB-FY24-0442.

Figure 2: Employment and Earnings Growth


(a) Employment Growth (1980–2022)



(b) Earnings Growth (1980–2022)




(c) Annual Wage Growth by Percentile




Notes: Data come from the 1980 Decennial Census and the 2010 and 2022 American Community Survey. Panels (a) and (b) show average annual employment growth and earnings growth, respectively. Nurse practitioners, CRNAs, and nurse midwives were not classified until the 2010 ACS. For that reason, the midlevel line in Panels (a) and (b) is for 2010–2022. Panel (c) shows annual wage growth at each percentile of the wage distribution. We focus on wages in this panel because measuring earnings at the very top of the earnings distribution using survey data is more difficult (Gottlieb et al., 2025). For context, we also print the 2022 level of wages at percentiles 10, 20, ..., 90, 99. For this panel, we restrict to respondents working at least 14 weeks a year for at least 20 hours on average to minimize measurement error when computing hourly wage, and we omit percentiles smaller than five because their baseline and final values are close to zero. For all plots, we calculate annual growth rates assuming continuous compounding. Monetary values are inflation-adjusted to 2022 dollars using the CPI-U. DRB number: CBDRB-FY24-0442.

Figure 3: Healthcare Employment: People

(a) Growth in Female Share (1980–2022)





(b) Growth in Foreign-Born Share (1980–2022)



Notes: This figure shows levels of, and changes in, the composition of healthcare employment by sex and foreign-born status. Data come from the 1980 Decennial Census and the 2010 and 2022 American Community Survey. The beginning of the arrow is the share in 1980 and the end is the share in 2022. Nurse practitioners, CRNAs, and nurse midwives were not classified until the 2010 ACS. For that reason, the midlevel line is for 2010–2022. The text to the right of each line is the annual percentage point change in the share. DRB number: CBDRB-FY24-0442.

) Change in Manufacturing Employment vs. Baseline Manufacturing Share



Appendix – For Online Publication

The Rise of Healthcare Jobs

Joshua D. Gottlieb, Neale Mahoney, Kevin Rinz, and Victoria Udalova

March 2025

A Data appendix

A.1 Industry Classification

In Figure 1, we use Current Employment Statistics (CES) from the BLS to plot industry employment over time. Specifically, we plot employment for the five largest major sectors by employment and the manufacturing major sector. We make two adjustments to the baseline CES data as defined by the BLS. First, we define healthcare by removing employees in social assistance from the BLS definition of “healthcare and social assistance.”²² Second, the baseline CES series is restricted to private-sector employees (and excludes public-sector employees). For this figure, we assign employees in public hospitals, public education, and public administrative roles to the healthcare, education, and “administrative, support, and waste management” series, respectively.

The rest of the paper uses the Census and the annual ACS. We code industries based on their 1990 definition and occupations based on their 2010 definitions—the latter since 2010 is the first year with midlevels. In the ACS, we use harmonized 1990 industry codes to define healthcare and manufacturing. We define healthcare as offices of physicians, dentists, chiropractors, optometrists, and other clinicians, hospitals, nursing and care facilities, and all other health facilities. Figure B.1 shows that this definition of healthcare yields employment estimates very close to the CES. We define manufacturing as all employment codes listed under the manufacturing group in the industry 1990 classification aside from printing and publishing industries and logging, which are not considered manufacturing in the CES and NAICS.

²²The social assistance subsector includes family services, food, housing, and relief services, vocational rehabilitation services, and child care services.

A.2 Occupation Classification

We use Census/ACS occupation codes to create our healthcare occupation categories. Because the harmonized coding scheme offered by IPUMS is based on 1990 codes, which do not break out nurse practitioners and CRNAs from nurses, we instead use codes contemporaneous to the survey year and harmonize accordingly.

Physicians are given a single code through 2018 until they are separated into physicians and surgeons. We include surgeons as physicians in those years. Nurses are defined as registered nurses (RNs) and licensed practical nurses (LPNs) throughout the sample. Until 2010, CRNAs and NPs were assigned the same code as nurses. Midlevels are defined as physician assistants, CRNAs, nurse practitioners, and nurse midwives; the latter two groups are combined by the ACS into a single category. Prior to 2010, the only distinct code in this group was physician assistants, so we omit statistics for midlevels until 2010.

Aides' codes become more granular over our sample. We use the broader grouping provided by IPUMS wherever possible. For example, in 1980, this includes the three occupation codes under "Health Service Occupations", while in 2017, this includes the 11 codes under "Health Support Occupations". The 2018–2022 codes do not include this header, so we select the codes that most closely match the ones in 2017.

"Other" clinical occupations is a catch-all category that includes all other patient-facing clinical occupation codes. For example, in 2018–2022, this includes most codes under "Healthcare Practitioners and Technical Occupations" that were not already covered by other occupation groups. There is considerable heterogeneity in this group, with dentists, physical therapists, and technicians all included. For this reason, we do not focus on this group on its own. "Non-clinical" occupations are defined as all remaining respondents with a healthcare industry code that are not matched into one of the previous groups. We cross-verify our crosswalk for aides, other clinical, and non-clinical using the 1990 codes from IPUMS to ensure we are capturing the same set of occupations each year. A complete list of codes can be found in Ruggles et al. (2024).

A.3 Restricted-Use Data

Our estimates using the Decennial Census or the ACS come from the restricted-use versions of the datasets. There are three advantages critical to our study. First, the internal files contain much larger samples. The public-use Decennial Census files include data from long-form surveys for roughly 5% of the US population. The restricted-use data contains long-form responses for 20% of the population. For ACS years, public-use files contain about two-thirds of all respondents for a given year. The internal files contain the full sample.

Second, income variables contain less income censoring (top-coding) in the internal versions. In the public-use files, income top-coding begins around the 99.5th percentile of the income distribution (McKenna and Haubach, 2019). An important part of our sample—physicians, in particular—has a large fraction of respondents censored in the public data. The restricted-use file has much less censoring. Gottlieb et al. (2023) use New York as an example to show that only 0.1% of the population and less than 0.5% of doctors have their incomes censored.


Third, we observe far more detailed geographic units in the restricted files. Most importantly, we observe counties, which allows us to conduct analysis using stable MSAs over time. Public-use files identify public-use microdata areas, which are no smaller than 100,000 people and are not consistent over time.

A.4 Calculating Healthcare’s Contribution to Female Employment

In Section 4.1, we note that 3.3 million of the 13.1 million women who found employment in 2022 due to rising female employment shares did so because of the growth in healthcare. We calculate 13.1 million by taking the difference between the share of women employed in 2022 and 1980 and multiplying by the working-age female population. We calculate 3.3 million by taking the difference in healthcare’s share of employment from 1980 to 2022 and multiplying by the working-age female population and the share of women employed in 2022.


B Additional exhibits

Figure B.1: Healthcare Employment Trends by Data Source



Notes: This figure shows the time series of employment for healthcare in the United States. Data come from the 1980, 1990, and 2000 Decennial Census, the 2001–2022 annual ACS survey, and the Current Employment Statistics (CES) from 1990–2022. We group ACS occupation codes to match BLS healthcare definitions. This iteration of the CES healthcare line based on NAICS codes is only available beginning in 1990. Appendix [A](#) contains further detail. DRB number: CBDRB-FY24-0442.


Figure B.2: Percentage Point Difference in Annual Wage Growth




Notes: This figure shows the difference in wage growth between specified industries from 1980 to 2022 at each percentile of the distribution in Figure 2c. Data come from the 1980 Decennial Census and the 2022 American Community Survey. To compute hourly wage, we take the self-reported annual wage and divide by the usual hours worked per week multiplied by the number of weeks worked. We restrict to respondents working at least 14 weeks a year for at least 20 hours on average to minimize measurement error when computing hourly wage. We omit the bottom five percentiles where average wages are close to zero and average annual growth is noisy. We calculate real annual growth rates assuming continuously compounding. The growth rate of a variable Y from its value Y_1 in year T_1 to Y_2 in year T_2 is $\frac{\ln(Y_2) - \ln(Y_1)}{(T_2 - T_1)}$. Wages are inflation-adjusted to 2022 dollars using the CPI-U. DRB number: CBDRB-FY24-0442.

Figure B.3: Average Annual Wage Growth by Percentile

(a) All employed respondents




(b) Respondents working at least 48 hours and 35 weeks




Notes: This figure shows the growth in wages by industry from 1980 to 2022 at each percentile of the wage distribution. Data come from the 1980 Decennial Census and the 2022 American Community Survey. To compute hourly wage, we take the self-reported annual wages and divide by the product of usual hours worked per week and the number of weeks worked. In Panel (a), we include all respondents. In Panel (b), we restrict to respondents working at least 48 weeks and at least 35 hours per week on average. We calculate real annual growth rates assuming continuously compounding. The growth rate of a variable Y from its value Y_1 in year T_1 to Y_2 in year T_2 is $\frac{(\ln(Y_2) - \ln(Y_1))}{(T_2 - T_1)}$. Wages are inflation-adjusted to 2022 dollars using the CPI-U. DRB number: CBDRB-FY24-0442.

Figure B.4: Change in Healthcare Share of Prime Age Population, MSAs (1980–2022)




Notes: This figure shows the density of the change in the healthcare share of the prime-age population from 1980–2022 for all MSAs and non-MSA state areas. Data come from the 1980 Decennial Census and the 2022 American Community Survey. The density curve is weighted by 2022 population. Prime-age population is defined as the total population of the geography between the ages of 25 and 54. DRB number: CBDRB-FY24-0442.

Figure B.5: Change in Non-Manufacturing, Non-healthcare Employment



Notes: This figure plots the change in the non-manufacturing, non-healthcare employed share of prime-age population against the manufacturing share of the prime-age population in 1980. Data come from the 1980 Decennial Census and the 2022 American Community Survey. “Non-manufacturing” and “Non-healthcare” refers to all employed respondents in an industry other than manufacturing or healthcare. Each dot represents a 2013 MSA or the parts of a state not in an MSA; only sufficiently large MSAs to satisfy Census Bureau release requirements are plotted. Prime-age population is defined as the total population of the geography between the ages of 25 and 54. DRB number: CBDRB-FY24-0442.


Figure B.6: 1980 Healthcare and 1980 Manufacturing Shares




Notes: This figure plots the 1980 manufacturing share of the prime-age population against the 1980 healthcare share of the prime-age population. Data come from the 1980 Decennial Census and the 2022 American Community Survey. Each dot represents a 2013 MSA or the parts of a state not in an MSA; only sufficiently large MSAs to satisfy Census Bureau release requirements are plotted. Prime-age population is defined as the total population of the geography between the ages of 25 and 54. DRB number: CBDRB-FY24-0442.

Figure B.7: Employment Across Space (MSAs)

(a) 2022 Healthcare Employment



(b) Change in Manufacturing Share, 1980–2022



Notes: Figure (a) maps the 2022 healthcare share of employment for MSAs. Figure (b) maps the change in the manufacturing share of prime-age population from 1980–2022 for MSAs. Data come from the 1980 Decennial Census and the 2022 American Community Survey. The 95 MSAs shown are sufficiently large to be disclosable under Census Bureau rules. We exclude Honolulu for visual purposes. DRB number: CBDRB-FY24-0442.

Table B.1: Earnings Growth, Education Decomposition

	Earnings Growth (1980-2022)		
	(1) Actual (%)	(2) Reweighted (%)	(3) Reweighted / Actual
Physicians	0.74	0.74	1.000
Nurses	1.29	1.09	0.843
Aides	0.83	0.75	0.907

Notes: Data from the 1980, 1990, and 2000 Census and the 2001–2022 ACS. We compute reweighted average annual growth following DiNardo, Fortin and Lemieux (1996) by reweighing the 2022 observations so that their education distribution matches the base period education distribution, and then estimating the growth rate. We calculate real annual growth rates assuming continuous compounding. The growth rate of a variable Y from its value Y_1 in year T_1 to Y_2 in year T_2 is $\frac{\ln(Y_2) - \ln(Y_1)}{(T_2 - T_1)}$. We omit midlevels because these occupations have only emerged recently, and their job roles and educational standards are still evolving. DRB number: CBDRB-FY24-0442.

Table B.2: Regression estimates, changing employment composition

	Δ in Share of Prime-Age Population					
	Disclosable MSAs			All MSAs		
	(1)	(2)	(3)	(4)	(5)	(6)
	Manufacturing	Healthcare	All Other	Manufacturing	Healthcare	All Other
Baseline Manufacturing Share	-0.60	0.07	0.60	-0.54	0.03	0.63
	(0.04)	(0.02)	(0.08)	(0.02)	(0.01)	(0.07)
Constant	0.02	0.03	0.05	0.02	0.04	0.05

Notes: This table reports coefficients from regressing the change in the share of the prime-age population employed in selected industries against the baseline manufacturing share. Data come from the 1980 Decennial Census and the 2022 American Community Survey. Prime-age population is defined as the total population of the geography between the ages of 25 and 54. The unit of observation is a 2013 MSA or the parts of a state not in an MSA. Disclosable MSAs satisfy Census Bureau release requirements. DRB number: CBDRB-FY24-0442.

Table B.3: Change in Employment and Baseline Manufacturing Employment

	1980 Share	2011 Share	Regression: emp. chg. vs. 1980 mfg. share	Regression: health emp. chg. vs. 1980 mfg. share Regression: mfg. emp. chg. vs. 1980 mfg. share
Panel A: All				
Manufacturing	0.160	0.080	-0.575	
Healthcare	0.056	0.084	0.058	0.100
Panel B: White Male (37%)				
Manufacturing	0.228	0.119	-0.748	
Healthcare	0.026	0.034	0.022	0.029
Panel C: White Female (37%)				
Manufacturing	0.095	0.046	-0.387	
Healthcare	0.078	0.130	0.094	0.243
Panel D: Nonwhite Male (13%)				
Manufacturing	0.206	0.095	-0.902	
Healthcare	0.039	0.044	-0.021	-0.023
Panel E: Nonwhite Female (14%)				
Manufacturing	0.114	0.048	-0.591	
Healthcare	0.115	0.135	0.052	0.088

Notes: This table relates the change in healthcare and manufacturing employment to a region's baseline manufacturing share. Data come from the 1980 Decennial Census and the 2011 American Community Survey. Healthcare and manufacturing are defined based on 1990 Census industry codes. We compute the 1980 and 2011 shares by taking the number of individuals ages 25–54 employed in manufacturing or healthcare and dividing by the total number of individuals ages 25–54. The third column reports the regression slope for the change in manufacturing or healthcare employment against a PUMA's baseline manufacturing share. The fourth column reports the ratio of the change in healthcare slope and the change in manufacturing slope. PUMAs refer to 1980–2011 “consistent” PUMAs defined by IPUMS.

Table B.4: Change in Employment and Baseline Manufacturing Employment

	1980 Share	2011 Share	Regression: emp. chg. vs. 1980 mfg. share	$\frac{\text{Regression: health. emp. chg. vs. 1980 mfg. share}}{\text{Regression: mfg. emp. chg. vs. 1980 mfg. share}}$
Panel A: All				
Manufacturing	0.160	0.080	-0.575	
Healthcare	0.056	0.084	0.058	0.100
Panel B: \leq Bachelor's Degree (70%)				
Manufacturing	0.172	0.084	-0.612	
Healthcare	0.050	0.073	0.054	0.089
Panel C: \geq Bachelor's Degree (30%)				
Manufacturing	0.110	0.070	-0.274	
Healthcare	0.080	0.110	0.059	0.215

Notes: This table relates the change in healthcare and manufacturing employment to a region's baseline manufacturing share. Data come from the 1980 Decennial Census and the 2011 American Community Survey. Healthcare and manufacturing are defined based on 1990 Census industry codes. We compute the 1980 and 2011 shares by taking the number of individuals ages 25–54 employed in manufacturing or healthcare and dividing by the total number of individuals ages 25–54. The third column reports the regression slope for the change in manufacturing or healthcare employment against a PUMA's baseline manufacturing share. The fourth column reports the ratio of the change in healthcare slope and the change in manufacturing slope. PUMAs refer to 1980–2011 “consistent” PUMAs defined by IPUMS.

C Theory appendix

We present a more detailed derivation of the theoretical result we discuss in Section 4.

C.1 Overview

Consider an economy with three industries: healthcare, manufacturing, and other. Let $s = \{H, M, O\}$ index those industries. We assume that each industry faces isoelastic exogenous demand

$$Q_s = Z_s P_s^\alpha \quad (\text{C.1})$$

where Q_s is the industry's output, P_s the unit price of the output, Z_s is a demand shifter, and $\alpha < 0$ the demand elasticity. Each industry has a Cobb-Douglas production function of the form

$$Q_s = A_s L_s^\beta K_s^{1-\beta} \quad (\text{C.2})$$

where L_s is labor, K_s is capital, and A_s is a productivity shifter. Total labor supply is inelastic at $\bar{L} = \sum_s L_s$ and workers are perfectly elastic across industries so equilibrium wages are constant and equal to the marginal revenue product of labor. Capital is also supplied perfectly elastically at a constant rental rate.

C.2 Cost minimization

With exogenous demand and perfect competition, each industry minimizes costs $wL + rK$ subject to equation (C.2). This yields a constant capital-to-labor expenditure ratio

$$\frac{K_s}{L_s} = \frac{w}{r} \frac{1-\beta}{\beta}. \quad (\text{C.3})$$

Substituting this into the production function (C.2) and solving for L_s in terms of Q yields

$$L_s = \frac{Q_s}{A_s} \beta^{1-\beta} (1-\beta)^{-(1-\beta)} w^{-(1-\beta)} r^{1-\beta}. \quad (\text{C.4})$$

We now solve for cost. Note that $wL + rK = wL \left[1 + \frac{rK}{wL}\right]$. Using (C.3), we can simplify and solve for cost

$$C(Q) = \frac{Q_s}{A_s} \beta^{-\beta} (1 - \beta)^{-(1-\beta)} w^\beta r^{1-\beta}. \quad (\text{C.5})$$

Perfect competition implies that the industry's output price is equal to the marginal cost $C'(Q)$:

$$P = \frac{1}{A_s} \beta^{-\beta} (1 - \beta)^{-(1-\beta)} w^\beta r^{1-\beta}. \quad (\text{C.6})$$

Substituting this into the demand equation (C.1) yields

$$Q_s = Z_s \left(\frac{1}{A_s} \beta^{-\beta} (1 - \beta)^{-(1-\beta)} w^\beta r^{1-\beta} \right)^\alpha. \quad (\text{C.7})$$

For simplicity, we collect constants by defining

$$\lambda = \beta^{-\beta} (1 - \beta)^{-(1-\beta)} r^{1-\beta} \quad (\text{C.8})$$

which allows us to express demand in terms of wages and productivity:

$$Q_s = Z_s \left(\frac{w^\beta \lambda}{A_s} \right)^\alpha. \quad (\text{C.9})$$

C.3 Labor demand and labor shares

To compute labor demand for industry s , we use the constant expenditure share $wL_s = \beta C(Q_s)$. We substitute our expression for the cost function, solve for L_s , and substitute for Q_s to obtain:

$$L_s = \frac{\beta \lambda w^{\beta-1} Q_s}{A_s} \quad (\text{C.10})$$

$$= \frac{\beta w^{\beta-1} \lambda}{A_s} \left[Z_s \left(\frac{w^\beta \lambda}{A_s} \right)^\alpha \right] \quad (\text{C.11})$$

$$= \beta w^{\beta(\alpha+1)-1} \lambda^{\alpha+1} \left[Z_s A_s^{-(\alpha+1)} \right]. \quad (\text{C.12})$$

Recall the adding-up constraint $\sum_s L_s = \bar{L}$. Healthcare's share of labor is thus

$$S_H = \frac{L_H}{\bar{L}} = \frac{Z_H A_H^{-(\alpha+1)}}{\sum_s Z_s A_s^{-(\alpha+1)}} \quad (\text{C.13})$$

C.4 Reallocation following demand shock

We ask how healthcare and “other” (non-healthcare, non-manufacturing) labor demand changes following a manufacturing shock.

Suppose there is a demand shock to manufacturing (via the demand shifter).²³ Then we have

$$\frac{\partial S_H}{\partial Z_M} = -\frac{A_M^{-(\alpha+1)} Z_H A_H^{-(\alpha+1)}}{\left(\sum_s Z_s A_s^{-(\alpha+1)}\right)^2} \quad (\text{C.14})$$

and

$$\frac{\partial S_O}{\partial Z_M} = -\frac{A_M^{-(\alpha+1)} Z_O A_O^{-(\alpha+1)}}{\left(\sum_s Z_s A_s^{-(\alpha+1)}\right)^2} \quad (\text{C.15})$$

A positive manufacturing demand shock decreases healthcare and other's shares of the labor force. To understand how these changes relate to baseline shares, we scale the changes in equations (C.14) and (C.15) by the industries' respective baseline shares:

$$\frac{\frac{\partial S_H}{\partial Z_M}}{S_H} = \frac{A_M^{-(\alpha+1)} Z_H A_H^{-(\alpha+1)}}{\left(\sum_s Z_s A_s^{-(\alpha+1)}\right)^2} \cdot \frac{\sum_s Z_s A_s^{-(\alpha+1)}}{Z_H A_H^{-(\alpha+1)}} = \frac{A_M^{-(\alpha+1)}}{\sum_s Z_s A_s^{-(\alpha+1)}} \quad (\text{C.16})$$

$$\frac{\frac{\partial S_O}{\partial Z_M}}{S_O} = \frac{A_M^{-(\alpha+1)} Z_O A_O^{-(\alpha+1)}}{\left(\sum_s Z_s A_s^{-(\alpha+1)}\right)^2} \cdot \frac{\sum_s Z_s A_s^{-(\alpha+1)}}{Z_O A_O^{-(\alpha+1)}} = \frac{A_M^{-(\alpha+1)}}{\sum_s Z_s A_s^{-(\alpha+1)}} \quad (\text{C.17})$$

These two quantities are identical. This means that each non-manufacturing industry has the same proportional response to a shock to manufacturing demand. This matches the results of our statistical analysis empirically connecting manufacturing job losses to healthcare job gains.

²³The conclusion is identical when there is a shock to manufacturing productivity.