Decomposing Historical Changes in Employer Costs and Contributions to Inequality*

Adam Bloomfield† Travis Cyronek‡ Christos A. Makridis§

February 10, 2025

Abstract

Using micro-data from the National Compensation Survey (NCS) across the U.S. private sector from 1990 to 2023, we introduce a methodology for decomposing changes in employer costs into price, hours worked, and employment distribution components across different industries, occupations, and firm sizes. Our findings reveal that the rising cost of benefits has been a significant contributor to the overall increase in total compensation, particularly in high-skill sectors such as Finance, Insurance, and Real Estate (FIRE). These rising costs have exacerbated wage polarization, particularly in low-skill services, by disproportionately benefiting higher-skilled workers. Our results are consistent with a growing valuation of non-wage amenities by employees.

Keywords: Structural Transformation; Employer Costs; Income Inequality; Wage Polarization; Compensation Decomposition; Sectoral Reallocation

JEL Codes: J31, J32, O41, D63, L16.

*The views expressed herein are those of the authors and do not necessarily reflect those of any affiliated institutions.

†Georgetown University Center for Retirement Initiatives, adam.bloomfield@gmail.com

‡Bureau of Labor Statistics, cyronek.travis@bls.gov

§University of Nicosia and Stanford University, cmakridi@stanford.edu

1 Introduction

The structure of labor costs in the U.S. has shifted dramatically over the past half-century with a higher proportion of total compensation coming from benefits over wages, growing from roughly 17% in the 1970s (Woodbury, 1983) to 30% by 2015 (Monaco and Pierce, 2015; Kristal et al., 2020). In the textbook model of a competitive labor market, wages primarily reflect the prevailing rate for a given job with little room for differentiation beyond skill level and industry. However, recent studies demonstrate that compensation is increasingly composed of both wages and non-wage benefits, which vary substantially by industry, occupation, and establishment size (Abowd et al., 1999; Sorkin, 2018). This shift has significant implications for how we understand employer costs, as benefits——particularly health care and retirement——have become a substantial and dynamic component of total compensation (Card et al., 2024). By contributing to wage polarization and reshaping industry wage structures, rising benefit costs highlight the importance of non-wage compensation in understanding the evolving labor market and income inequality (Pierce, 2001).

Understanding the determinants of *compensation* has become increasingly relevant as non-wage benefits reshape the structure of employer costs and influence workers' preferences. Compensation packages differ widely not only across industries, but also across occupations and firm sizes, creating a complex landscape of wage and benefit distribution (Sorkin, 2018; Card et al., 2013). High-wage sectors and large firms, for example, often provide more generous benefits, while lower-wage industries may offer fewer amenities, contributing to variations in compensation that go beyond base pay. This trend is not solely an artifact of market power; it also reflects firms' strategic use of benefits to attract and retain talent, particularly in sectors where working conditions are demanding or competitive (Lamadon et al., 2022; Sockin, 2022). As benefits increasingly complement wages in shaping compensation, they play a central role in defining job quality and, ultimately, worker welfare, adding new dimensions to the study of income inequality and labor market dynamics.

The primary contribution of this paper is to study the evolution of the costs to employ labor, distinguishing between not only price versus quantity effects, but also wage versus non-wage compensation across different industries, occupations, and establishment sizes. Using the National Compensation Survey (NCS), fielded by the Bureau of Labor Statistics (BLS), we introduce a methodological approach to decompose cumulative changes in employer costs stemming from changing prices, hours worked, and the distribution of employment. First, aggregate trends show steady growth in real compensation since 1990 with benefits increasing more rapidly than wages. This divergence accelerated during the early 2000s, slowed slightly

after the 2008 financial crisis, and experienced a decline after 2020. The key driver of this growth is rising prices, particularly for benefits like health insurance and retirement contributions, as reflected in the sharp increases in the benefits price sub-index. Second, benefits have become an increasingly significant component of compensation, outpacing wages across industries, occupations, and establishment sizes. This trend is particularly pronounced in high-wage sectors such as Finance, Insurance, and Real Estate, high-skill occupations like Management and Professional roles, and smaller establishments, where benefits have contributed disproportionately to overall compensation growth. High-wage earners have also seen the largest increases in benefits, further widening inequality. Third, the hours worked and employment composition have remained relatively stable, contributing minimally to the overall growth in compensation costs. Slight declines in hours worked were observed, especially in high-benefit-cost establishments, while shifts in employment distribution toward higher-cost industries, occupations, and larger establishments modestly increased costs, particularly after 2013. Fourth, there is significant heterogeneity in compensation growth by industry and occupation. High-wage sectors and high-skill occupations show the most substantial increases in both wages and benefits, driven by price effects, while lower-wage sectors and lower-skill roles exhibit slower growth. This pattern highlights the role of skill-based differentials in shaping compensation dynamics. Fifth, smaller establishments face disproportionately high growth in benefits costs compared to larger firms, reflecting their limited bargaining power and reduced economies of scale. In contrast, larger firms demonstrate more stable compensation growth with steady, but moderate, benefits growth. Finally, we use our created index to estimate the effects of changes in the price of benefits, holding fixed changes in composition and quantity, on inequality, accounting for X% since 1990.

The first part of our paper introduces the compensation data of the NCS and describes several measurement hurdles that may obfuscate the comparison of estimates across time. We discuss the changing way establishments and workers are classified and the potential for bias stemming from a rotating sample, ultimately producing an index we take to the data and decompose the price effect, hours effect, and employment effect. Whereas the Laspeyres price effect reflects the change in costs using baseline period quantities, the Paasche employment effect reflects the changes in employment due to cost-by-the-hour. The hours effect is a hybrid of the two, capturing the changing work schedule using cost-by-employment weights from different periods. We produce these indices for total compensation, benefits, and wages.

The second part of our paper uses this methodology to study the changing landscape of compensation, as well as both wages and non-wage benefits, in private establishments from 1990 to 2023 across industry, occupations, and establishment sizes. We document three

empirical regularities. First, benefits are increasingly outpacing wages across all dimensions of the labor market with especially steep growth in high-wage industries (like Finance and Real Estate), high-skill occupations (such as Management and Professional roles), and smaller establishments. Furthermore, non-wage benefits are greater among higher wage jobs, suggesting complementarity between the two. Second, price changes are the main driver behind these cost increases, with a stronger impact on benefits than on wages. Across all industries, occupations, and establishment sizes, rising benefit prices—rather than shifts in hours worked or employment composition—account for the majority of cost growth. Smaller establishments, in particular, may feel the weight of these rising prices due to their limited bargaining power and reduced economies of scale. Third, growth in wages and benefits vary by group characteristics. High-wage sectors, high-skill roles, and smaller establishments experience faster growth in both compensation components, but with a sharper increase in benefits. Lower-wage sectors and lower-skill roles see more modest growth in both wages and benefits, though benefits still contribute more heavily to cost increases.

Following these empirical regularities, we use our new indices to understand whether the rise in the price of non-wage benefits can help explain the labor market polarization that has taken place over the past few decades. A historical analysis of compensation trends across the wage distribution highlights a U-shaped pattern: compensation growth is strongest at the bottom and top of the distribution, but weaker in the middle. While wage growth drives compensation increases for lower-wage jobs, benefit growth is the primary driver at the upper end. This trend has intensified since 2000. Price changes, rather than shifts in hours worked or employment composition, are the dominant force behind rising benefit costs. Growth in benefit prices consistently outpaces wage price growth, particularly for high-wage jobs. Employers appear to adjust to these rising costs by shifting employment patterns, reducing hours for higher-paid positions while increasing hours for lower-wage jobs. Overall, these findings indicate that benefits are a growing driver of inequality in total compensation with high-wage workers receiving a larger share of their earnings in non-wage forms.

1.1 Related Literature

Our paper contributes most to the literature on income inequality and polarization in wage growth by linking these phenomena to the decomposition of employer costs. The rapid pace of technological change has led to increased wage polarization, particularly among low-skilled workers (Autor, 2006; Autor et al., 2008; Autor and Dorn, 2013). We extend this discussion by examining how the rising costs of benefits and other non-wage components of compensation contribute to this polarization. For instance, we show that the sharp

increase in benefit costs in sectors like FIRE has exacerbated income inequality by disproportionately rewarding high-skilled workers while squeezing low-wage sectors. Furthermore, by decomposing employer costs, we shed light on the differential impacts of these costs on various occupational groups, highlighting how the increasing burden of benefit costs is contributing to broader trends of wage stagnation and inequality. Our findings suggest that these (de-)compositional dynamics are likely to intensify income polarization, especially as technological advancements continue to reshape the labor market (Acemoglu and Restrepo, 2022).

Ever since at least Woodbury (1983), there has been a recognition that benefits are a meaningful portion of total compensation, growing to roughly 17% by the 1970s. Pierce (2001) subsequently showed that standard measures of wages understate the degree of inequality due to unobserved differences in non-wage benefits that are correlated with skill.[1] Motivated by the continued rise of benefits as a share of compensation to roughly 30% (Monaco and Pierce, 2015), some have argued that the expansion of benefits has led to an increase of inequality within organizations, benefiting "core employees" in higher paid roles more (Kristal et al., 2020). Others have examined how these benefits are a tool that organizations use to attract and positively select specific types of employees (Liu et al., 2023).

Our study also relates to the empirical microeconomics literature concerning the causes and consequences of employer choices to offer non-wage compensation to workers. In a seminal study, Freeman (1981) uses establishment data to show that unionism increases the share of total compensation devoted to fringe benefits. A series of studies in the mid-1980s found that income tax rates, and rate changes, have a statistically significant but small effect on employer expenditures on non-wage benefits, the fringe share of compensation, and the number of fringe benefits offered (Long and Scott, 1982, 1984; Turner, 1987b,a). Using the National Medical Care Expenditure Survey (NMCES), Kiker and Rhine (1987) finds that significant determinants of total fringe benefits are worker educational level, age, marginal tax rate, female sex, full-time employment, and white collar employment. Oyer (2008) develops a model that incorporates scale economies in benefit costs, search costs for workers who value non-wage compensation, and disutility from work and using NLSY data finds empirical support for all three factors. More recently, using data from Glassdoor, Liu

[1]Performance related pay has also been attributed as one such factor (Lemieux et al., 2009), but others find that it plays an insignificant role (Gittleman and Pierce, 2015). That is not to say that performance related pay is not important, but rather that it could effect earnings through selection margins at the establishment-level (Makridis and Gittleman, 2022). Performance pay could also have incentive effects on individuals to accumulate more human capital, which may not show up as much at the establishment-level if there are other unobserved determinants of human capital (Makridis, 2022).

et al. (2023) show that firms use maternity leave to attract female employees in labor markets where human capital, in the form skilled female workers, is relatively scarce. Furthermore, Makridis (2019) shows that employees are willing to pay for non-pecuniary amenities, such as feeling appreciated and having better relationships with their manager, and Makridis and Schloetzer (2024) show that these factors dwarfed the role of remote work as a non-pecuniary amenity during the pandemic. Our decomposition of benefits allows us to isolate the role of price effects over other long-run shifts, such as the composition of jobs and/or quantity effect, complementing emerging evidence on the value employees place on non-wage benefits.

Prior research has found that that the cost of non-wage compensation, specifically health care benefits, is offset by reductions in workers' wages (Summers, 1989; Gruber and Krueger, 1991; Baicker and Chandra, 2006). And recent scholarship in health economics has found that increasing cost benefits caused by price growth for health services can function as a regressive "head tax". These relationships can theoretically result in relatively larger negative employment effects among lower-wage workers relative to higher-wage employees (Finkelstein et al., 2023). However, Brot-Goldberg et al. (2024) recently find that rising health care benefit costs have no measurable disemployment and wage growth affects among the lowers income workers, who seldom receive health insurance or other nonwage compensation from their employers (Lurie and Miller, 2023).

2 Measurement Exercise

In this section we introduce our main source of data, the National Compensation Survey (NCS), and develop an index number that allows us to decompose changes in compensation—and its components—arising from shifts in the price of compensation, the volume of hours worked, and then structure of representative worker (the distribution of employment across industry, occupation, and establishment size). Important to this work is the consideration and treatment of several measurement hurdles that may obfuscate the comparison of data collected across distant samples. In particular we contend with variability that arises from a rotating sample and the aggregation of observations.

2.1 National Compensation Survey

Our main source of data come from the NCS, an establishment survey with a rotating panel structure that collects information on the costs to employ labor. It considers a broad definition of compensation: in addition to wages, data are collected on the costs to provide several types of benefits, including paid leave, supplemental pay, insurance, retirement and

savings, and some legally required plans. The survey is conducted quarterly, covering US employment in the Private Sector and State and Local Governments.[2] Unique establishment-occupations form *jobs*, which serve as the principal unit of observation. Costs, for both wages and benefits, are collected as averages across all workers within a given job.[3] Hours data are also collected to convert these costs to an hourly rate ("per hour worked"). Importantly, the hours measures are collected at an establishment's initiation into the survey and are updated only if there has been a permanent change to the work schedule. If the employees in a job work fewer hours in a particular reference week, but are still expected to work the same number of hours over the course of the year, the hours variable is not updated.[4]

Though the name *NCS* has been used since 1996, collection of the compensation data that would eventually become it dates back into the late 1970s (Bureau of Labor Statistics, 2017). To facilitate the comparison of these data over such long periods of time, we must ensure the samples represent comparable populations. This is handled by adjusting sampling weights to match the distribution of employment across certain dimensions. A challenge arises due to changes in the way establishments, and jobs, are classified over time. Industries in early samples were defined by the Standard Industrial Classification (SIC) system, before adopting the North American Industry Classification System (NAICS). Occupations, before the Standard Occupational Classification (SOC), were categorized according to the Occupational Classification System (OCS). Using crosswalks between these systems and a Multinomial Logit model to assign occupation groups for some earlier samples (which do not have detailed codes), we adjust sample weights to target the distribution of national employment across industry, occupation, and establishment size (see Table A.1 for the groups). Using the Quarterly Census of Employment and Wages (QCEW) for industry-by-size counts and the Current Population Survey (CPS) for occupation counts, we construct final weights for Private Sector workers from 1990 to 2023 (see Appendix C for further details on weighting).

2.2 Index Number

We start with a ratio of *value aggregates* in arbitrary periods 1 (the measuring period) and 0 (the base period). Letting $j \in J$ and $\widetilde{j} \in \widetilde{J}$ index the sets of jobs in 1 and 0, respectively:

[2]The NCS excludes federal workers, military personnel, agricultural workers, volunteers, unpaid workers, the self-employed, and those employed in private households (Bureau of Labor Statistics, 2017).

[3]We, therefore, do not observe costs for individual workers.

[4]In other words, this collection strategy does not pick up seasonal fluctuations.

$$R_1^{total} = \frac{\sum\limits_{j \in J} c_{j1} h_{j1} w_{j1}}{\sum\limits_{\widetilde{j} \in \widetilde{J}} c_{\widetilde{j}0} h_{\widetilde{j}0} w_{\widetilde{j}0}}. \tag{1}$$

c is the price of labor (as a cost per hour worked) and is either *compensation*, *wages*, or *benefits*.[5] Each component c is made *real* using CPI-U by Census Region. h is the work schedule (the average number of weekly hours worked), and w is the weight (which reflects the distribution of employment across industry, occupation, and establishment size). Since the values for c and h are averages across all workers in a job, the *total relative* presented in Equation 1 captures the total growth in real weekly average labor costs between the measuring and base periods. Since Equation 1 may be sensitive to outliers, we trim the bottom and top 0.5% of records based on their job-level growths for wages and benefits.

Next, we point out that sampling variability between the sets J and $\widetilde{J}$ may obfuscate the measurement of change between periods. For example, if the sample in 1 happens to capture higher cost establishments than in 0 (beyond what can be controlled for by the weights), Equation 1 will be upward biased. To address this issue we exploit the panel dimension of the NCS and restrict the sets J and $\widetilde{J}$ to be equal. That is, we require the same jobs be present in both measuring and base periods, and drop all indexing of jobs $j \in J$ for notational brevity.

To decompose changes in compensation, take Equation 1 and multiply by "1" *twice*, rearranging terms to yield the following:

$$R_1^{total} = \underbrace{\frac{\sum c_1 h_0 w_0}{\sum c_0 h_0 w_0}}_{R_1^{price}} \times \underbrace{\frac{\sum c_1 h_1 w_0}{\sum c_1 h_0 w_0}}_{R_1^{hours}} \times \underbrace{\frac{\sum c_1 h_1 w_1}{\sum c_1 h_1 w_0}}_{R_1^{emp}}. \tag{2}$$

The expression in Equation 2 separates the total growth into a pure price effect (the price relative, R_1^{price}), an hours effect (the hours relative, R_1^{hours}), and an employment effect (the employment relative, R_1^{emp}). R_1^{price} is Laspeyres-like in that the change in costs uses base-period quantities (where hours-adjusted employment are the quantities). R_1^{emp}, in contrast, is Paasche-like in that changes due to employment are cost-by-hour weighted using values in

[5] Compensation = Wages + Benefits.

the measuring period. Finally, R_1^{hours} is a hybrid: it measures the changes due to a changing work schedule using cost-by-employment weights from different periods. A nice property of Equation 2 is that it is multiplicatively separable, a feature that will carry over to the final index number calculation.

The final step of this exercise is to accumulate these relatives to permit assessments of change across time intervals of arbitrary length. In the early NCS, establishments are sampled for up to 5 years, but only up to 3 years in later samples. The requirement of matched-observations in the base and measuring periods limits the use of these relatives to at most 5 or 3 years. Further, because of the rotation of sample members, measurements across longer intervals will not utilize all collected records. Some records collected in 1 might not have been present when collection in 0 occurred (i.e. newly rotated sample members), and some that were collected in 0 were not in 1 because they rotated out. To accommodate these challenges, we adopt a *chained* approach. We calculate quarterly changes, letting the base and measuring periods be $t-1$ and t, respectively. This maximizes the sample size for any given interval. To measure changes over longer periods of time, we accumulate these relatives, linking them by multiplying consecutive periods together. Our final index number, measured between periods 0 and T, is given by

$$I_T^{total} = \prod_{t=0}^{T} R_t^{total} = \underbrace{\left(\prod_{t=0}^{T} R_t^{price}\right)}_{I_T^{price}} \underbrace{\left(\prod_{t=0}^{T} R_t^{hours}\right)}_{I_T^{hours}} \underbrace{\left(\prod_{t=0}^{T} R_t^{emp}\right)}_{I_T^{emp}}. \tag{3}$$

Equation 3 provides for a total index number, multiplicatively separable into three sub-indexes for prices, hours, and employment. By construction $I_0 = 1$ for each index and sub-index. For $t > 0$ each index can be used to evaluate the change in real weekly average labor costs arising from changes in prices, the work schedule, the distribution of employment, or all of the above. The above can also be applied to different cuts of data. Records are conditioned on industry, occupation, establishment size, etc., to provide disaggregated index series for comparison. Combined with our efforts calibrating sample weights to a target distribution of employment that is comparable over distant periods of time makes this a novel exercise measuring continuous changes in real compensation costs using the NCS.

3 Results

3.1 Aggregate Changes in Compensation, Wages, and Benefits

Figure 1 plots our private establishment index series for compensation, wages, and benefit costs from 1990 to 2023. The top-left panel reports the *total* index, while the other three panels separately report our decomposition due to changing prices, the volume of hours, and the structure of the representative worker (the distribution of employment across occupation, industry, and establishment size). We discuss now discuss these patterns.[6]

The top-left panel shows the total index for compensation, wages, and benefits. It measures the cumulative change in real weekly costs to employ the representative worker. It reflects the combined effects of changes in the prices of labor services, the volume of hours worked, and the structure of the employment distribution. The compensation index (blue line) shows a steady increase from 1990 onward, with a more rapid rise during the early to mid-2000s. The growth slows slightly after 2008, likely reflecting the financial crisis, but then resumes its upward trend that eventually accelerates until around 2020, after which there is a decline. The wage index (orange line) also shows an increase over the period, with a similar pattern to total compensation but at a slightly slower rate. Again, the early to mid-2000s period exhibits the fastest growth for this index. The benefits index (green line) shows the most dramatic increase, growing more rapidly than wages, particularly after the early 2000s and disproportionately contributing to the overall rise in real weekly compensation costs.

The top-right panel reports the price sub-index, which measures cumulative changes in real weekly costs to employ the representative worker stemming only from the changing prices of labor services (i.e. keeping the volume of hours and structure of the employment distribution fixed). The compensation price sub-index shows a strong upward trend, particularly during the early to mid-2000s, which is the period of fastest growth. After 2008, the growth rate stabilizes, continuing to rise steadily until around 2020. The wage price sub-index shows a similar upward trend, with the most rapid increase occurring in the early 2000s. This line follows the general trend of compensation but at a slightly lower rate. The benefits price sub-index shows the steepest rise among the three components, particularly in the early 2000s. The increasing price of benefits, such as health insurance and retirement contributions, has been a major driver of the overall rise in compensation costs.

The bottom-left panel shows the hours sub-index, which measures cumulative changes

[6] Table 1 in the Online Appendix summarizes these growth rates in a tabular format. Price changes are the main driver of growth in compensation, wages, and benefits, particularly in recent decades with benefits seeing the highest increases.

in real weekly costs to employ the representative worker stemming only from the changing volume of hours worked (i.e. keeping the prices of labor services and the employment distribution fixed). The compensation hours sub-index shows a slight decline over the period, indicating a gradual reduction in average weekly hours worked. The early 1990s see some fluctuations, but the trend stabilizes by the mid-2000s, with a slight downward drift. The wages hours sub-index mirrors total compensation the compensation hours sub-index, showing this small decline, especially before 2008. In contrast, the decline observed in the benefits hours sub-index is greater, reflecting that hours are decreasing more for relatively high benefit cost establishments. That said, since wages are a larger fraction of compensation, both the wages and compensation hours sub-indexes are closer to each other.

The bottom-right panel shows the employment distribution sub-index, which measures cumulative changes in real weekly costs to employ the representative worker stemming from only changes to the distribution of employment across industry, occupation, and establishment size (i.e. keeping the prices of labor services and the volume of hours fixed). The compensation employment sub-index shows a modest declining trend through much of the period before starting to rise around 2013. The sub-index rises sharply around 2020, reflecting the significant shifts in employment during that period. The wage employment sub-index follows a similar pattern to compensation, with an even more modest decline in the early period. The benefits employment sub-index shows a steeper decline through 2011 prior to increasing more sharply than the other two sub-index lines. When these sub-indexes are declining, we can infer that employment is being shifted away from higher cost industries, occupations, and establishment sizes. Conversely, when these sub-indexes increase, employment is shifting into higher cost industries, occupations, and establishment sizes.

3.2 Heterogeneity

3.2.1 By Industry

Figure 3 studies compensation trends across different industries. In the compensation panel (top left), sectors like Finance, Insurance, and Real Estate and Services exhibit strong growth, reflecting higher wage and benefit increases in these areas. The price sub-index shows that price changes have driven much of this compensation growth with Finance and Real Estate seeing notable price-related increases. The hours and employment sub-indexes indicate that hours worked remain relatively stable across industries, whereas employment distribution shifts slightly, revealing sectoral adjustments over time.

Next, we disentangle the role of benefits and wage compensation. Figure 4 starts with the

wage trends by industry. Consistent with the compensation results, the wage (top left) panel shows industries such as Finance and Construction experiencing significant wage growth. The price sub-index reflects strong price-driven increases in Finance, whereas Manufacturing and Utilities maintain more stable trends. The hours and employment sub-indexes demonstrate that hours remain relatively unchanged, while sectoral shifts in employment composition contribute marginally to wage cost variations.

Finally, Figure 5 shows the benefit cost trends across industries. Again, there are rapid increases across all sectors, particularly in Finance and Real Estate. The Total Benefits panel shows consistent growth, while the price sub-index reveals the relative price sensitivity of benefits in certain sectors. The Hours Sub-Index shows relatively minor adjustments, and the employment distribution sub-index suggests gradual shifts in the composition of employment within different industries, implying an evolving labor structure that impacts benefit costs.

We conclude with three broad observations. First, benefit costs have risen more sharply than wage costs in most industries, especially in high-wage sectors like Finance, Insurance, and Real Estate. Second, price changes are the main reason for both wage and benefit growth, but the impact is stronger for benefits. In industries like Finance, the rise in benefit prices has been more substantial than in wages. Third, Since 2000, benefits have consistently outpaced wages in growth, particularly in high-cost sectors. This highlights a trend where benefits are increasingly important in compensation packages across industries.

3.2.2 By Occupation

Figure 6 documents compensation trends across occupational groups. Management, Professional, and Related occupations have the highest compensation growth, likely due to increased benefits and wages for high-skill roles. The price sub-index supports this with robust growth across occupational groups, particularly in high-skill categories. The hours sub-index remains stable across occupations, while the employment sub-index indicates gradual structural adjustments in occupational composition over time.

Figure 7 focuses on wage trends across occupations. The Total Wages panel highlights increased wages primarily in higher-skill roles. The price sub-index shows steady increases across all groups, with Management and Professional roles leading. The hours and employment sub-indexes show limited change, indicating that most wage growth is driven by price adjustments rather than hours worked or changes in occupational composition.

Figure 8 turns towards benefit costs. The total index (top left) indicates significant growth in benefit costs across all occupations, with the most considerable increases observed in high-skill occupations, such as Management, Professional, and Related. The price sub-index highlights that price changes are the main drivers of these increases, while the hours and employment sub-indexes remain relatively stable.

We conclude with three broad observations. First, and like before, benefit costs have increased more rapidly than wage costs across all occupation groups. High-skill occupations, such as Management, Professional, and Related. Second, price changes are the main driver of cost growth. However, price growth is more pronounced for benefits than for wages in each occupational category. Management and Professional occupations see larger cumulative increases in both wages and benefits compared to other roles, but the difference is sharper for benefits. Service and Sales and Office occupations show slower wage and benefit growth, with a more modest increase in benefits relative to high-skill roles. Third, for high-skill roles (e.g., Management and Professional), the difference in growth between benefits and wages highlights a shift towards benefits-heavy compensation. This trend is less pronounced in lower-skill roles where wage and benefit growth are more balanced but still slower.

3.2.3 By Establishment Size

Figure 9 documents compensation trends by establishment size. The total index (top left) shows that larger establishments (500+ employees) have had more stable compensation growth compared to smaller firms, which experienced more pronounced fluctuations. The price sub-index reveals that price changes drive compensation increases across all establishment sizes, with smaller establishments facing higher cost growth. Both the hours and employment sub-indexes indicate minimal changes, pointing to price factors as the primary source of compensation cost variation across establishment sizes.

Figure 10 displays wage trends segmented by establishment size. The total index suggests that smaller establishments experienced faster wage growth than larger ones. The price sub-index shows that wage increases were primarily driven by price factors, particularly in smaller firms. The hours and employment sub-indexes remain relatively constant across establishment sizes, reinforcing the conclusion that price factors, rather than structural changes, are responsible for the observed wage growth.

Figure 11 shows benefits costs by establishment size. The total index demonstrates significant growth in benefits across all establishment sizes, with smaller establishments experiencing more rapid increases. The price sub-index highlights that benefits cost growth

is predominantly driven by price changes, particularly in smaller firms. The hours and employment sub-indexes show minimal variation, suggesting that the rising benefits cost burden is largely a function of price increases across all establishment sizes.

We conclude with three broad observations. First, benefits have risen more rapidly than wages, especially among smaller establishments (fewer than 50 workers). This suggests that smaller firms may be shifting a larger portion of compensation growth into benefits, potentially to attract and retain employees without substantial wage hikes. Second, price increases are the main driver of growth in both wages and benefits, but they have a greater impact on benefits costs, particularly in smaller firms. This indicates that smaller establishments face a heavier financial burden from rising benefit prices, possibly due to less bargaining power and fewer economies of scale compared to larger firms. Third, hours worked and employment composition have remained relatively stable across establishment sizes, indicating that changes in labor costs are almost entirely due to price factors.

3.3 Complementarity of Wage and Non-Wage Benefits

For the last part of our analysis, we study the relationships between the growths of wage and non-wage costs. Figure 2 compares average wage levels with benefit cost growth for different cuts of the labor market. Occupational groups with higher average wages tend to have higher benefit growth, indicating that high-paying roles also tend to incur higher benefit costs over time. Conversely, lower-wage occupation groups like Service exhibit slower growth in benefits. For industry groups, Finance, Insurance, and Real Estate stands out for having both high wages and significant benefit growth. For establishment sizes, the largest establishments–with the highest wages–also have the highest growth in benefit costs. This visualization suggests a positive correlation between wage levels and benefit cost growth, particularly in high-wage industries, occupations, and establishment sizes.

Next, we revisit some earlier analysis on inequality to explore these relationships across the wage distribution (Pierce, 2001, 2010; Monaco and Pierce, 2015). The cited work utilizes NCS microdata and estimates various relationships between compensation costs for different wage percentiles. Our analysis differs a few ways. First, due to the availability of more data, our period of study extends through the 2010s. Second, our sample weights are constructed specifically for comparability of distant samples. The earlier work utilized weights similar to those from the Employer Costs for Employee Compensation (ECEC) data product, published by the BLS, which do not adjust for changes in occupational group shares and have not always made adjustments for changes to the distribution of establishment sizes. Further, there have

been changes to how the ECEC weights have been constructed over time, notably due to the adoption of NAICS to classify industries. In contrast, our weighting procedure is designed to produce the distribution of employment across industry, occupation, and establishment size across industry and occupation classification regimes (SIC and NAICS for industry, OCS and SOC for occupation).[7],[8] Third, and last, we utilize the index number apparatus developed above to measure changes in real costs. Pierce (2010), for example, uses log differences at different points in time to measure changes. Instead, our index estimates will measure cumulative changes and, importantly, allow us to decompose these changes in terms of changing prices, the volume of hours, and the distribution of employment.

Unlike the above analysis, which uses static characteristics to define the domain of an estimate, estimation for a specific percentile may be impossible because there could be no qualifying observations. We therefore follow the approach employed by Pierce (2010). First, we sort all jobs by their average wage and use hours-adjusted sample weights to identify percentile cutoffs. Next, we bin observations using a 20pp window centered on each percentile and then compute an estimate. That is, for the 50th percentile, we identify all records between the 40th and 60th percentile before making computations. For the 51st percentile, we utilize records between the 41st and 61st percentiles, etc. This moving window strategy smooths the percentile estimates in a non-parametric way so that jobs with relatively large frequency weights do not produce big jumps when calculating weighted percentiles. [9]

Figure 12 reports the fraction of compensation in benefits by wage percentile at different points in time. In 2020, benefit costs accounted for about a fifth of compensation costs at the 10th percentile of wages, but about a third near the 75th percentile. This falls slightly at the top of the distribution where very high wages dilute the share of benefits. Second, benefits have become a larger share of compensation for all wages, but especially higher wage jobs. Growth was slow, if not slightly negative, from 1990 to 2000. However, the first decade of the millennium saw substantial growth above the 20th percentile of wages. This trend continued into the 2010s for jobs above the 30th percentile. Taken together, the share of compensation in benefits has grown and has grown more for higher wage jobs.

In Figure 13 we calculate and plot this paper's total index values across the wage distribu-

[7]For details on how this is accomplished, see Appendix C.

[8]In addition to what is mentioned here, the Pierce work utilizes an hours-adjustment to its weights for comparability with CPS sample weights. Our index formulation specifically incorporates this hours adjustment and separates it out for decomposition purposes.

[9]Note also that our smoothing is more aggressive than the original work. This is in anticipation of variance concerns associated with our chained index estimator. Having too few observations can produce noisy estimates that make calculations susceptible to chain drift.

tion for the 1990–2020 period and for each intervening decade. In keeping with the chained formulation of our index, we re-bin records for every quarterly change and then accumulate these changes over time. For the entire period of study we find a familiar U-shape highlighted by prior work: compensation (including wages and benefits) growth is high near the bottom of the wage distribution, falls for the middle of the distribution, and is high again for the upper end of the distribution. When compensation growth (blue line) is above wage growth (orange line), benefit growth is relatively high. The distance between these lines suggests how large this benefit growth is. We can therefore see that the growth in compensation at the bottom of the wage distribution is driven by growth in wages, while growth in compensation at the top of the wage distribution is driven by growth in benefits. In the other panels of Figure 13 we break this down by decade. The 1990s have a growth profile most emblematic of this U-shape, with a small amount of separation between compensation and wage growth throughout. This separation increases substantially in the 2000s, indicating that benefit growth is dominant. In addition, benefit growth is higher the further right one moves across the wage distribution. Finally, in the 2010s growth is much flatter for the highest wage jobs, and high growth for the lowest wage jobs is driven by wages, not benefits.

Last, we decompose the changes for 1990–2020 into the sub-components allowed by our index construction. Figure 14 documents the total and sub-index values for compensation costs (including wages and benefits) for different wage percentiles. The first panel makes clear that benefit growth increases in wage percentile. Though there is some compression in compensation at the very bottom of the distribution, benefits are positively contributing to inequality. In the top right panel, we can see that most of this is comes from prices: the price of wages increases modestly across the distribution, but is much lower than the growth in benefit prices. In the bottom left panel we see that hours have a negative but minor effect for most of the distribution, but are important for compensation growth at the very bottom of the wage distribution. This indicates that establishments, on the margin, are purchasing more hours for the lowest paid jobs and slightly fewer hours for the highest paid jobs. In the bottom right panel are the employment sub-indexes. Recall that this sub-index measures the cumulative change in real costs arising *only from* movement in employment across industry, occupation, and establishment size (holding prices and hours fixed). This allows us to see how employers may be responding to price increases. From the plot, as one moves from the left to the right, we see a declining growth until the 80th percentile of wages, where growth is more neutral. This pattern is consistent with the idea that, as labor gets more expensive, producers substitute away from those types of jobs on the margin. That is, until they cannot: jobs in the right tail of wages may be less substitutable.

4 Conclusion

The structure of labor costs in the United States has shifted over the past several decades with non-wage benefits playing an increasingly central role in total compensation. Using the National Compensation Survey (NCS) from the Bureau of Labor Statistics (BLS), we introduce a strategy for decomposing changes in employer costs into price, hours, and employment distribution effects to better understand the drivers of compensation growth.

Three key stylized facts emerge from our analysis. First, benefit costs have outpaced wage growth across all dimensions of the labor market, particularly in high-wage sectors (such as Finance and Real Estate), high-skill occupations (like Management and Professional roles), and smaller establishments.[10] This reflects an increasing reliance on non-wage compensation, especially in roles and sectors with traditionally high earnings. Second, price increases are the dominant driver of rising compensation costs with benefit price growth exceeding that of wages. Changes in benefit costs are primarily driven by rising prices—rather than shifts in hours worked or workforce composition—placing a heavier financial burden on employers, especially in benefits-heavy compensation structures. Third, compensation trends vary by group characteristics, with high-wage sectors, high-skill occupations, and smaller firms experiencing steeper increases in both wages and benefits, though the growth in benefits is particularly pronounced. Meanwhile, lower-wage sectors and lower-skill occupations exhibit more modest increases, highlighting how industry, skill level, and firm size shape compensation structures and cost growth over time.

Several questions remain for future research. First, how do rising benefit costs influence firm hiring and wage-setting behavior, particularly for lower-wage workers who receive fewer non-wage benefits? Second, given that benefits costs are largely driven by price increases, what policy mechanisms, such as health care reforms or changes to tax incentives, could mitigate their regressive effects on labor costs? Third, to what extent do rising benefits shape worker mobility, job satisfaction, and long-term employment outcomes? Lastly, as remote

[10] Table A.2 in the Online Appendix documents compensation cost growth by price changes, showing high growth in sectors like Finance and Real Estate and in smaller establishments, which lack the economies of scale of larger firms. Across occupational groups, costs rise steadily, driven largely by the price of benefits, which aligns with overall trends indicating rising non-wage costs. Table A.3 in the Online Appendix conducts an analogous exercise, but for hours worked. While hours reductions are slightly more prominent in larger establishments, overall there is minimal variation in hours-driven cost growth, suggesting that other factors, such as price changes, play a larger role in driving total compensation changes. Similarly, Table A.4 in the Online Appendix shows real growth in costs due to employment distribution shifts. Larger firms show negative growth, indicating a shift toward smaller or medium-sized establishments. The occupational adjustments show stable or positive growth in high-skill roles, while lower-skill roles experience downward adjustments, reflecting employment redistribution toward higher-value-added occupations.

work and gig employment expand, how will the composition of compensation continue to evolve, and what implications will this have for labor market inequality and job quality? Addressing these questions will help refine our understanding of compensation structures and their broader economic consequences.

References

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.

Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in u.s. wage inequality. *Econometrica*, 90(5):1973–2016, 2022.

David H Autor. The polarization of the u.s. labor market. *American Economic Review*, 96 (2):189–194, 2006.

David H Autor and David Dorn. The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–1597, 2013.

David H. Autor, Lawrence F. Katz, and Melissa S. Kearney. Trends in u.s. wage inequality: Revising the revisionists. *Review of Economics and Statistics*, 90(3):300–323, 2008.

Katherine Baicker and Amitabh Chandra. The labor market effects of rising health insurance premiums. *Journal of Labor Economics*, 24(3):609–634, 2006.

Zarek Brot-Goldberg, Zack Cooper, Stuart V Craig, Lev R Klarnet, Ithai Lurie, and Corbin L Miller. Who pays for rising health care prices? evidence from hospital mergers. Technical report, National Bureau of Economic Research, 2024.

Bureau of Labor Statistics. National compensation measures. *Handbook of Methods*, 2017. URL https://www.bls.gov/opub/hom/ncs/home.htm.

David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.

David Card, Jesse Rothstein, and Moises Yi. Industry wage differentials: A firm-based approach. *Journal of Labor Economics*, 42(S1):S11–S59, 2024.

Amy Finkelstein, Casey McQuillan, Owen Zidar, and Eric Zwick. The health wedge and labor market inequality. *Brookings Papers on Economic Activity*, 2023(1):425–503, 2023.

Richard B. Freeman. The effect of unionism on fringe benefits. *ILR Review*, 34(4):489–509, 1981.

Maury Gittleman and Brooks Pierce. Pay for performance and compensation inequality: Evidence from the ecec. *Industrial and Labor Relations Review*, 68(1):28–52, 2015.

Jonathan Gruber and Alan B Krueger. The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance. *Tax policy and the economy*, 5: 111–143, 1991.

B. F. Kiker and Sherrie L. W. Rhine. Fringe benefits and the earnings equation: A test of the consistency hypothesis. *The Journal of Human Resources*, 22(1):126–137, Winter 1987.

Tali Kristal, Yinon Cohen, and Edo Navot. Workplace compensation practices and the rise in benefit inequality. *American Sociological Review*, 85(2):271–297, April 2020.

Thibaut Lamadon, Magne Mogstad, and Bradley Setzler. Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212, 2022.

Thomas Lemieux, W. Bentley MacLeod, and Daniel Parent. Performance pay and wage inequality. *Quarterly Journal of Economics*, 124(1):1–41, 2009.

Tim Liu, Christos A. Makridis, Paige Ouimet, and Elena Simintzi. The distribution of nonwage benefits: Maternity benefits and gender diversity. *Review of Financial Studies*, 36 (1):194–234, 2023.

James E. Long and Frank A. Scott. The income tax and nonwage compensation. *Review of Economics and Statistics*, 64(2):211–219, May 1982.

James E. Long and Jr. Scott, Frank A. The impact of the 1981 tax act on fringe benefits and federal tax revenues. *National Tax Journal*, 37(2):185–194, June 1984.

Ithai Z Lurie and Corbin L Miller. Employer-sponsored health insurance premiums and income in us tax data. *Journal of Public Economics*, 224:104942, 2023.

Christos A. Makridis. Does culture pay? evidence from crowdsourced employee engagement data. Technical report, SSRN, 2019. Working Paper.

Christos A. Makridis. Work and grow rich. *SSRN working paper*, 2022.

Christos A. Makridis and Maury Gittleman. On the cyclicality of real wages and employment: New evidence and stylized facts from performance pay and fixed wage jobs. *Journal of Law, Economics, and Organization*, 38(3):889–920, 2022.

Christos A. Makridis and Jason Schloetzer. Does working from home increase job satisfaction and retention? evidence from the covid-19 pandemic. *Management Science, R&R*, 2024.

Kristen Monaco and Brooks Pierce. Compensation inequality: Evidence from the national compensation survey. *Monthly Labor Review*, 2015.

Paul Oyer. The making of an investment banker: Stock market shocks, career choice, and lifetime income. *Journal of Finance*, 63(6):2601–2628, December 2008.

Brooks Pierce. Compensation inequality. *Quarterly Journal of Economics*, 116(4):1493–1525, 2001.

Brooks Pierce. Recent trends in compensation inequality. In *Labor in the new economy*, pages 63–98. University of Chicago Press, 2010.

Jonathan Sockin. Show me the amenity: Are higher-paying firms better all around? Working paper, SSRN, 2022.

Isaac Sorkin. Ranking firms using revealed preference. *The Quarterly Journal of Economics*, 133(3):1331–1393, 2018.

Lawrence H Summers. Some simple economics of mandated benefits. *The American Economic Review*, 79(2):177–183, 1989.

Robert W. Turner. Are taxes responsible for the growth in fringe benefits? *National Tax Journal*, 40(2):June 1987, June 1987a.

Robert W. Turner. Taxes and the number of fringe benefits received. *Journal of Public Economics*, 33(1):41–57, June 1987b.

Stephen Woodbury. Substitution between wage and nonwage benefits. *American Economic Review*, 73(1):166–182, 1983.

Appendix

A Tables

Component	1990–2000	2000 – 2010	2010 – 2020
		ESTIMATES	
Compensation			
Prices	7.2	7.5	6.2
Hours	-0.1	-0.4	-0.4
Employment	-0.7	-1.2	3.7
Total	6.3	5.9	9.6
Wages			
Prices	5.9	3.7	6.0
Hours	0.1	-0.3	-0.4
Employment	-0.2	-0.4	3.6
Total	5.9	3.0	9.4
Benefits			
Prices	10.6	18.1	6.7
Hours	-0.5	-0.7	-0.5
Employment	-2.2	-3.1	3.9
Total	7.7	13.7	10.4
		STANDARD ERRORS	
Compensation			
Prices	0.5	0.5	0.3
Hours	0.7	0.4	0.1
Employment	0.1	0.1	0.1
Total	0.9	0.6	0.3
Wages			
Prices	0.6	0.5	0.3
Hours	0.7	0.4	0.1
Employment	0.1	0.1	0.1
Total	1.0	0.6	0.4
Benefits			
Prices	0.8	0.8	0.4
Hours	1.0	0.3	0.1
Employment	0.1	0.1	0.1
Total	1.4	0.9	0.4

Table 1: *Growth Rates by Decade.* This table reports real growth rates by decade of compensation components. Using our index number formulation, we decompose these growth rates in terms of changing prices, hours worked, and the distribution of employment across industry, occupation, and establishment size. For example, total compensation increased by 7.2% in the 1990s due to changing prices to employers (holding hours worked and the distribution of employment fixed). Total compensation decreased by 0.1% in the 1990s due to reduced hours (holding prices and the employment distribution fixed). Finally, it decreased by 0.7% due to shifts in employment (holding prices and hours fixed). Calculations are performed with corrections for the total sum of weights so that "Employment" growth reflects shifts in the distribution of employment and *not* population growth. Last, we use a chained index approach where weight shares are updated each quarter and changes are accumulated over time.

B Figures

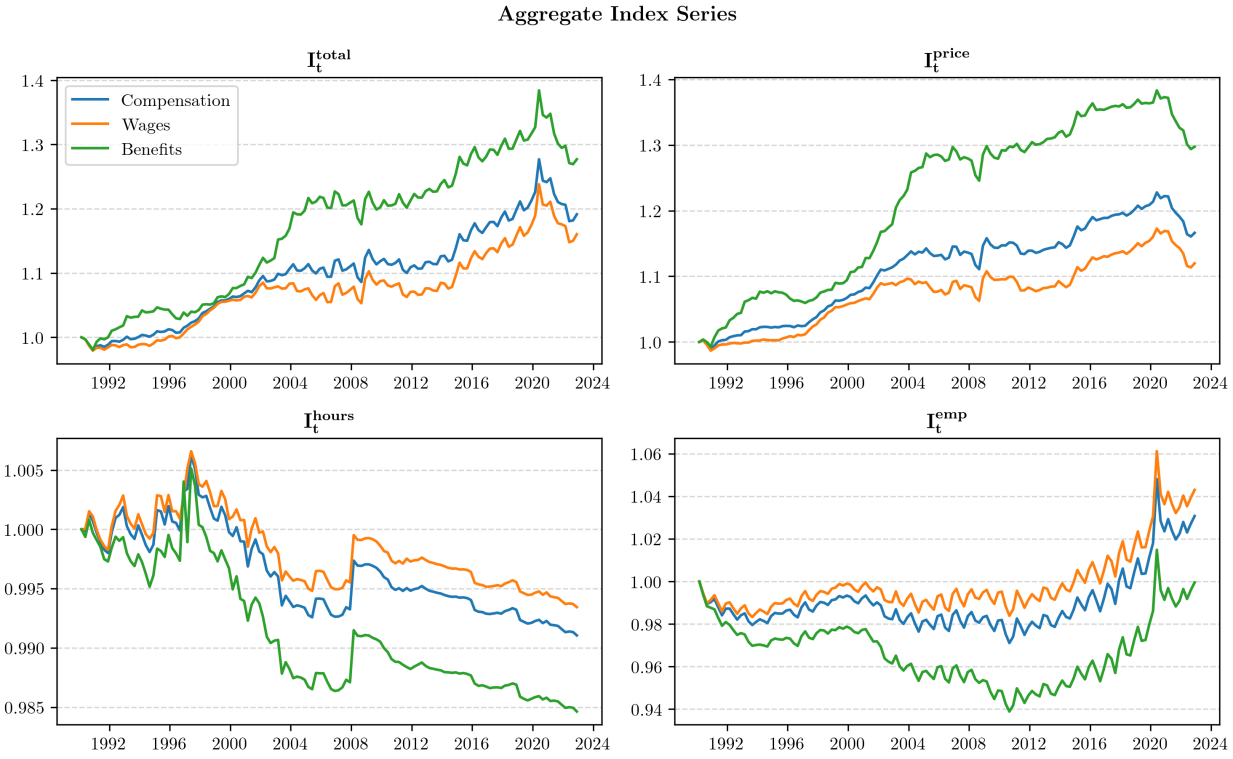


Figure 1: *Index Series.* The above figure plots our index series of compensation, wages, and benefit costs for Private establishments starting in the first quarter of 1990. The top-left panel reports changes in the value aggregates ("total"), while the other three panels separately report our decomposition. The top right panel is the price index, while the bottom left is a volume index for hours worked, and the bottom right is an index for the changing structure of the employment distribution. The product of these three decomposed series equals the top-left series.

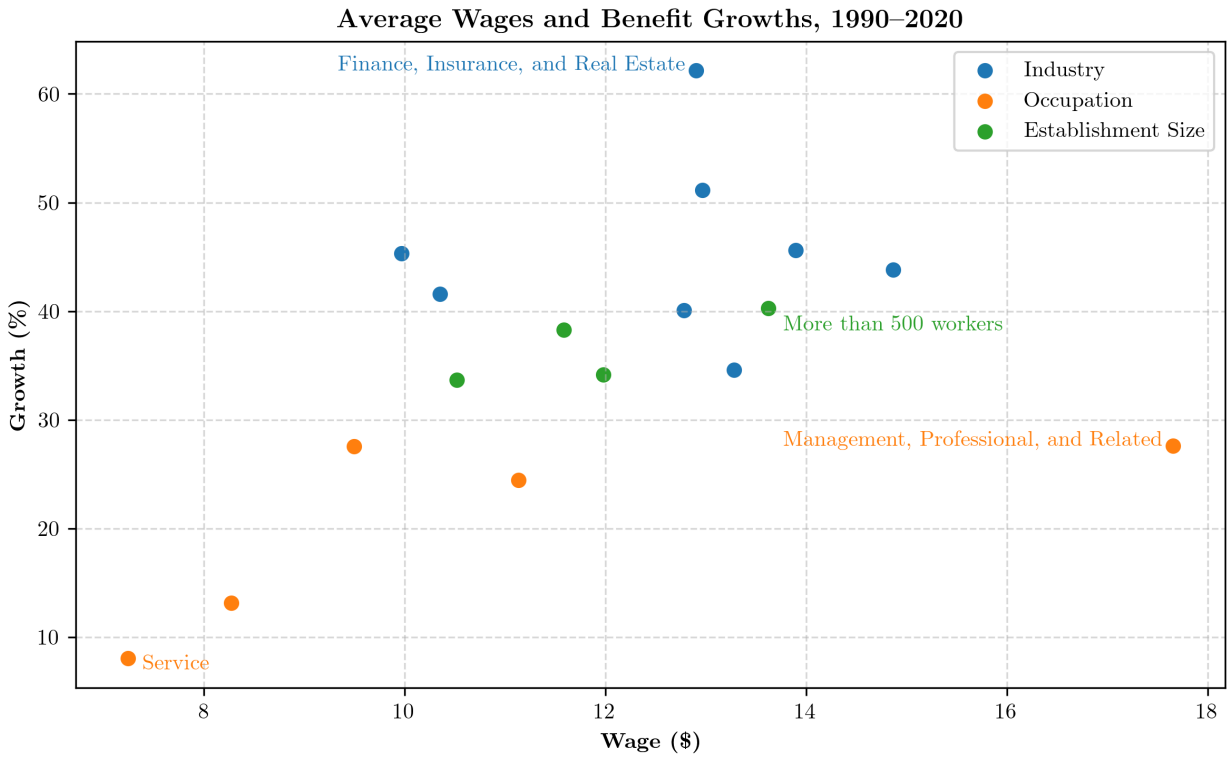


Figure 2: *Average Wages and Benefit Growths by Group, 1990–2020 (annotated).* The above figure plots a scatter of average wages (in 1990 dollars) by the total real growth in benefit costs (I^{total} for benefits) for different domains (industry, occupation, and establishment size) from 1990 to 2020.

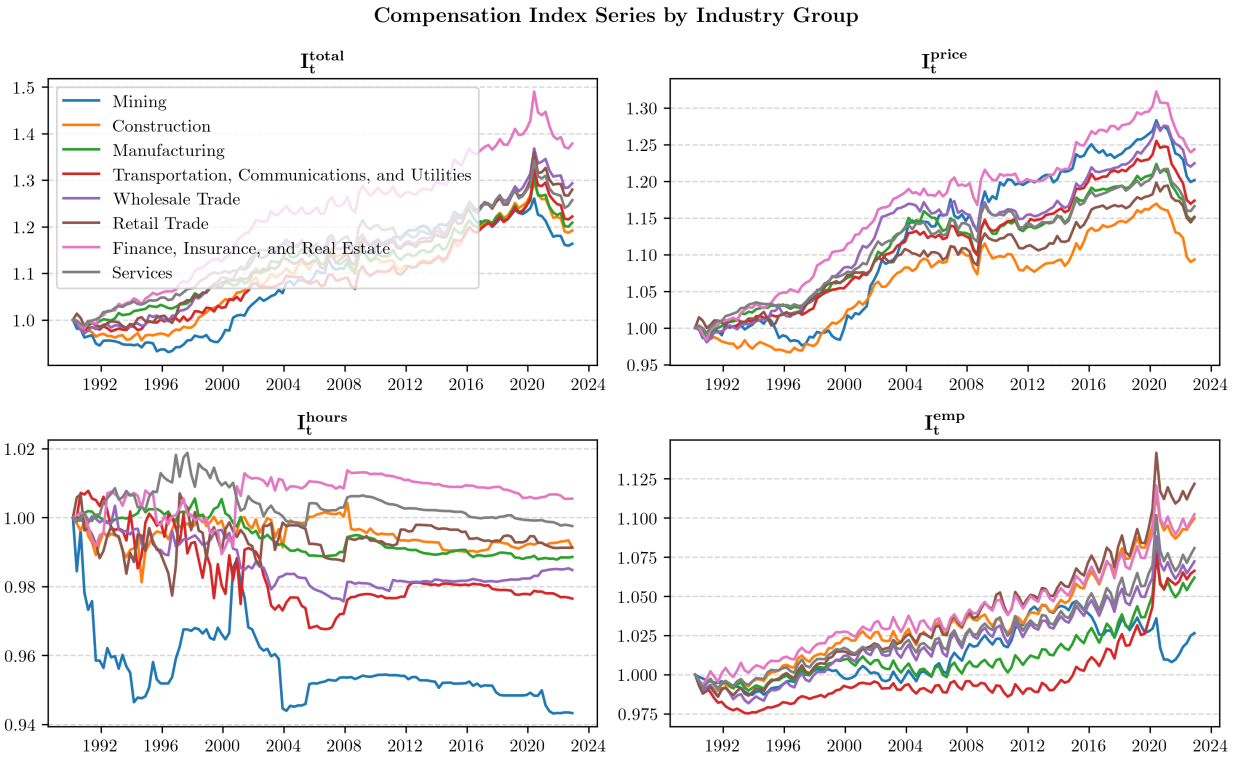


Figure 3: *Compensation Index Series by Industry Group.* The above figure plots our index series of compensation costs broken down by industry group starting in the first quarter of 1990. The top left panel reports the total change in compensation costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various industry groups, the employment distribution indexes should be interpreted as the change compensation costs from shifts in employment across the occupation and establishment size dimensions (that is, holding the industry distribution fixed).

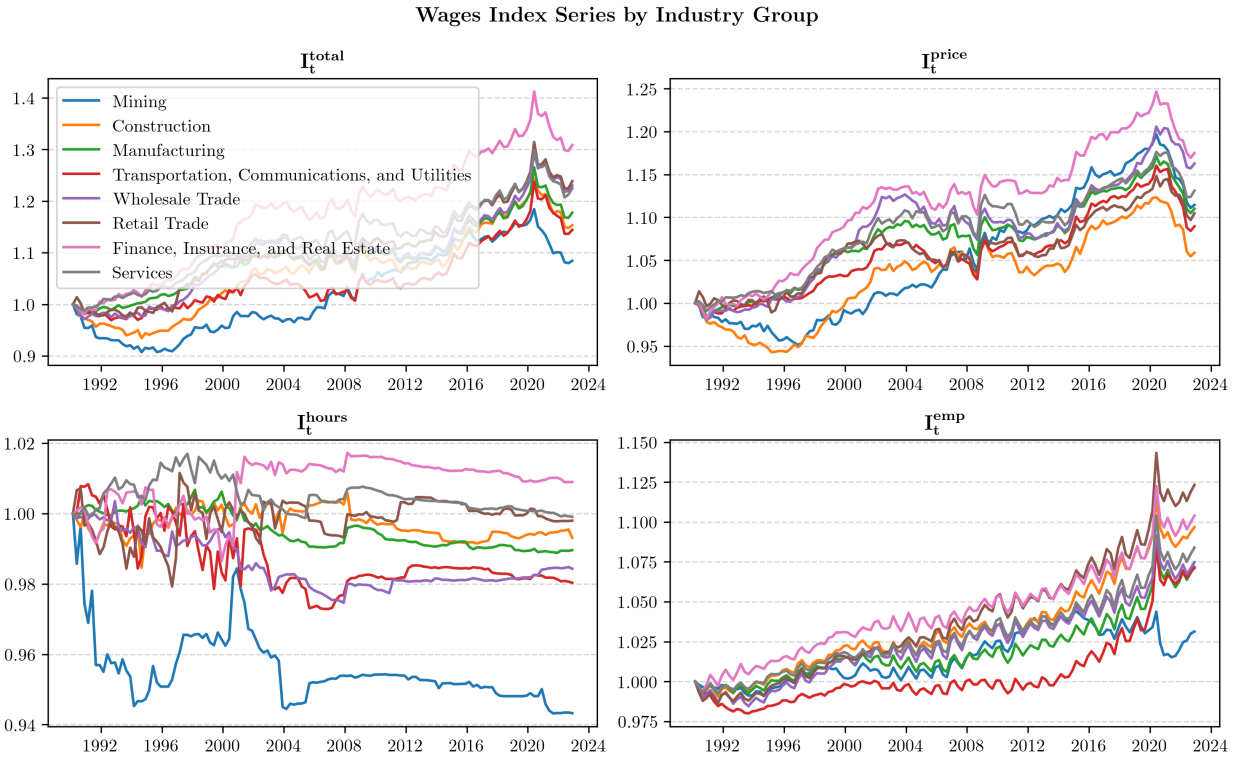


Figure 4: *Wage Index Series by Industry Group.* The above figure plots our index series of wage costs broken down by industry group starting in the first quarter of 1990. The top left panel reports the total change in wage costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various industry groups, the employment distribution indexes should be interpreted as the change wage costs from shifts in employment across the occupation and establishment size dimensions (that is, holding the industry distribution fixed).

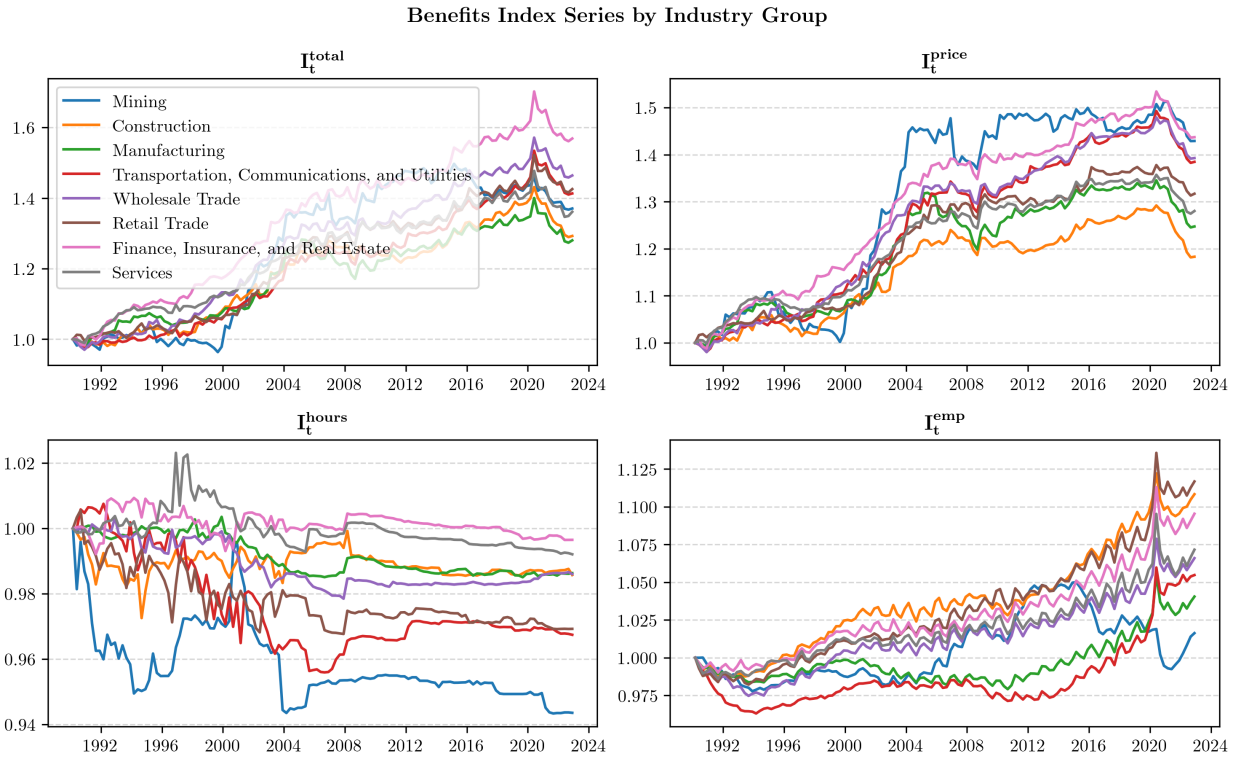


Figure 5: *Benefits Index Series by Industry Group.* The above figure plots our index series of benefit costs broken down by industry group starting in the first quarter of 1990. The top left panel reports the total change in benefit costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various industry groups, the employment distribution indexes should be interpreted as the change benefit costs from shifts in employment across the occupation and establishment size dimensions (that is, holding the industry distribution fixed).

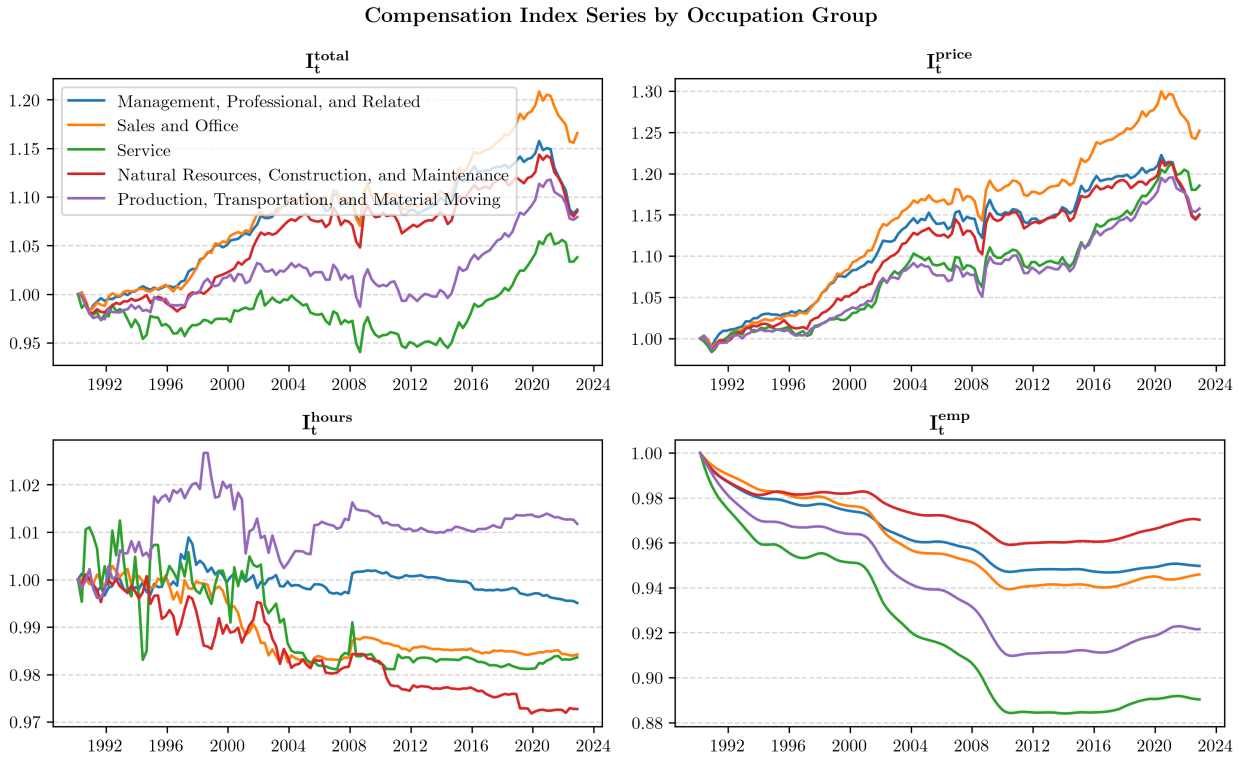


Figure 6: *Compensation Index Series by Occupation Group.* The above figure plots our index series of compensation costs broken down by occupation group starting in the first quarter of 1990. The top left panel reports the total change in compensation costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various occupation groups, the employment distribution indexes should be interpreted as the change compensation costs from shifts in employment across the industry and establishment size dimensions (that is, holding the occupation distribution fixed).

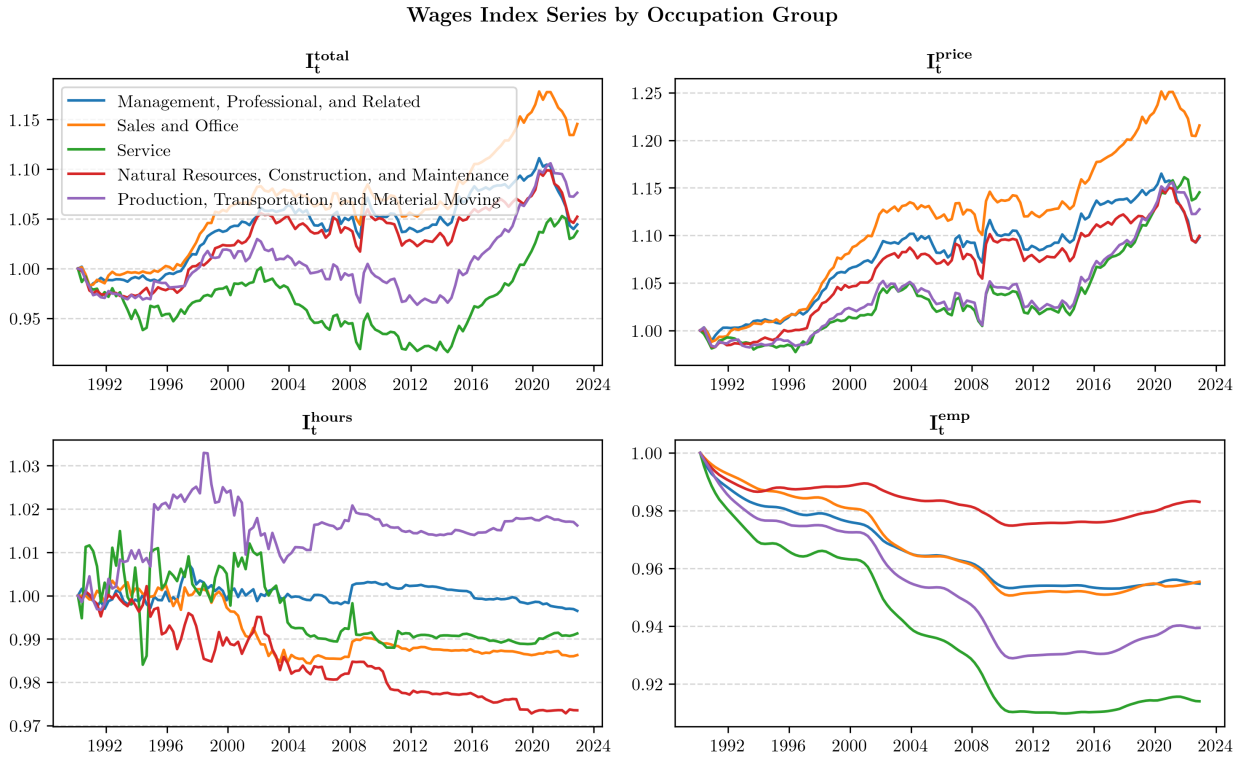


Figure 7: *Wage Index Series by Occupation Group.* The above figure plots our index series of wage costs broken down by occupation group starting in the first quarter of 1990. The top left panel reports the total change in wage costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various occupation groups, the employment distribution indexes should be interpreted as the change wage costs from shifts in employment across the industry and establishment size dimensions (that is, holding the occupation distribution fixed).

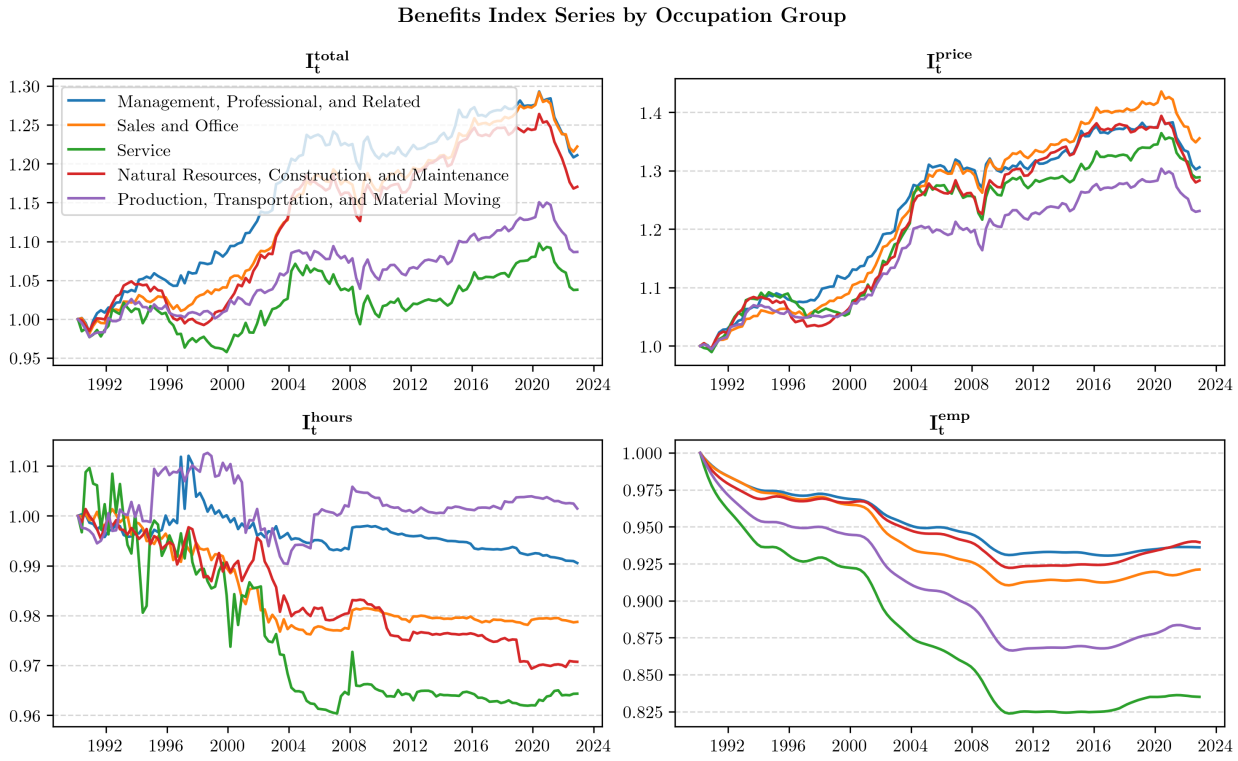


Figure 8: *Benefits Index Series by Occupation Group.* The above figure plots our index series of benefit costs broken down by occupation group starting in the first quarter of 1990. The top left panel reports the total change in benefit costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various occupation groups, the employment distribution indexes should be interpreted as the change benefit costs from shifts in employment across the industry and establishment size dimensions (that is, holding the occupation distribution fixed).

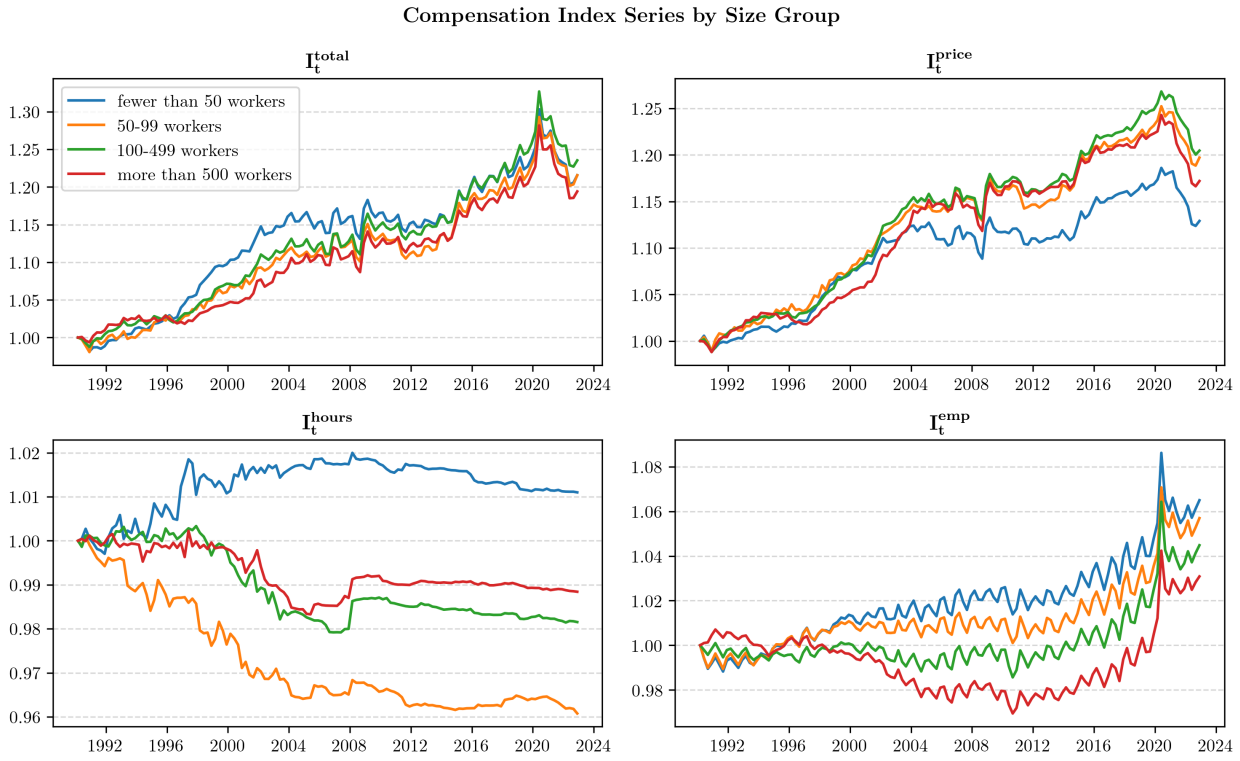


Figure 9: *Compensation Index Series by Establishment Size Group.* The above figure plots our index series of compensation costs broken down by establishment size group starting in the first quarter of 1990. The top left panel reports the total change in compensation costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various establishment size groups, the employment distribution indexes should be interpreted as the change compensation costs from shifts in employment across the industry and occupation dimensions (that is, holding the establishment distribution fixed).

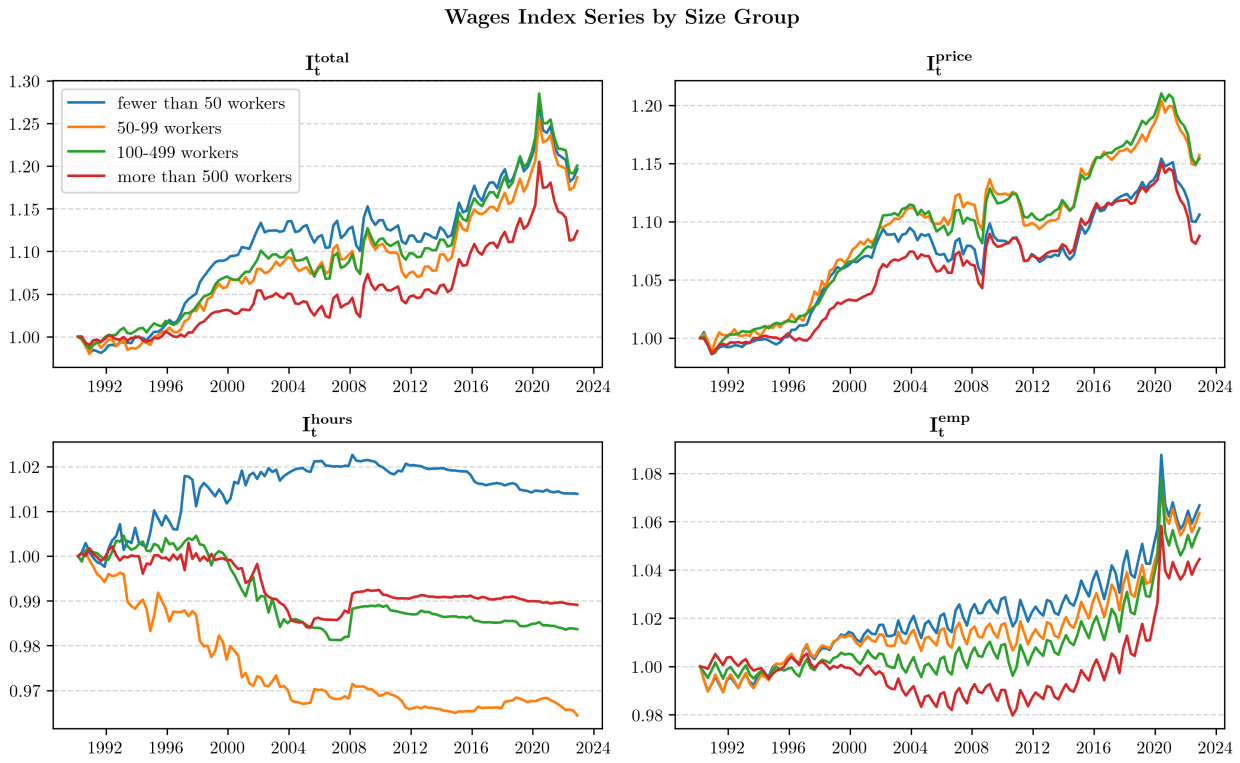


Figure 10: *Wage Index Series by Establishment Size Group.* The above figure plots our index series of wage costs broken down by establishment size group starting in the first quarter of 1990. The top left panel reports the total change in wage costs, while the other three panels report the price, hours, and employment distribution decomposition. Since these indexes condition on the various establishment size groups, the employment distribution indexes should be interpreted as the change wage costs from shifts in employment across the industry and occupation dimensions (that is, holding the establishment distribution fixed).

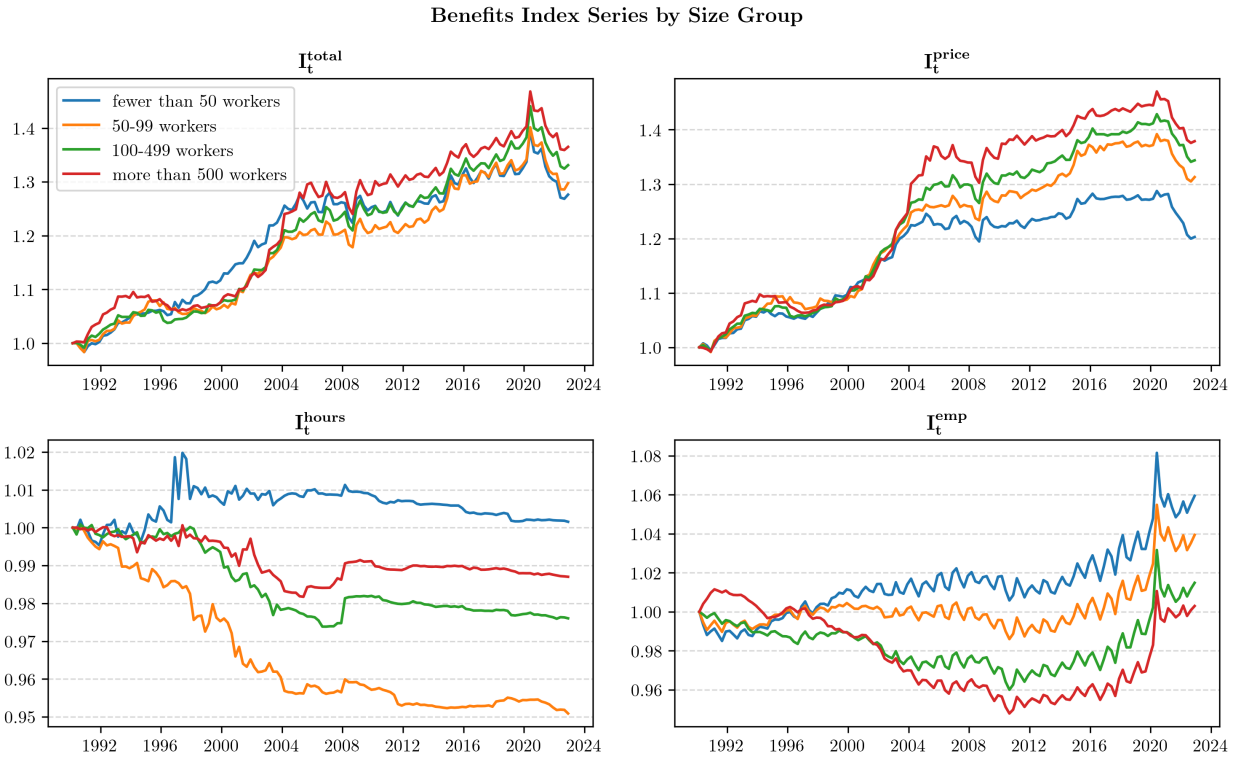


Figure 11: *Benefits Index Series by Establishment Group.* The above figure plots our index series of benefit costs broken down by establishment size group starting in the first quarter of 1990. The top left panel reports the total change in benefit costs, while the other three panels report the price, hours, and employment distribution indexes. Since these indexes condition on the various establishment size groups, the employment distribution indexes should be interpreted as the change benefit costs from shifts in employment across the industry and occupation dimensions (that is, holding the establishment distribution fixed).

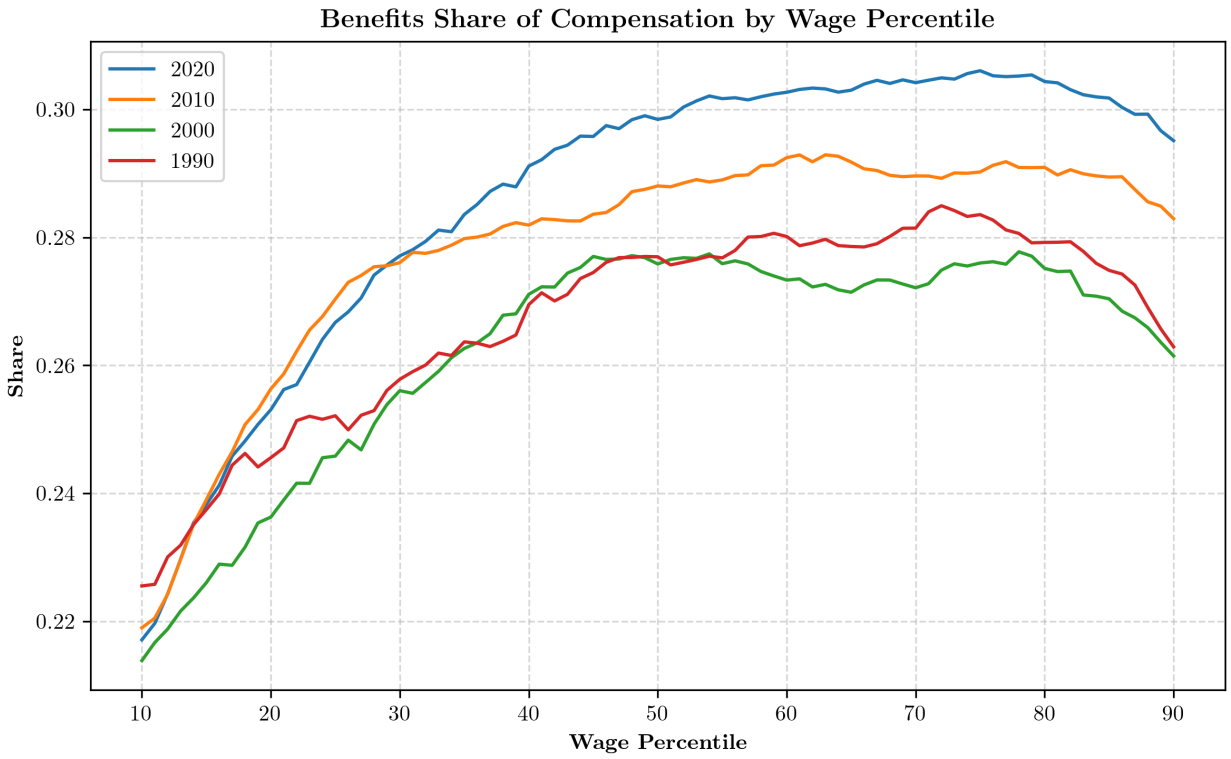


Figure 12: *Fraction of Compensation in Benefits by Wage Percentile* The above figure plots the share of benefits in compensation across the wage distribution for 1990, 2000, 2010 and 2020.

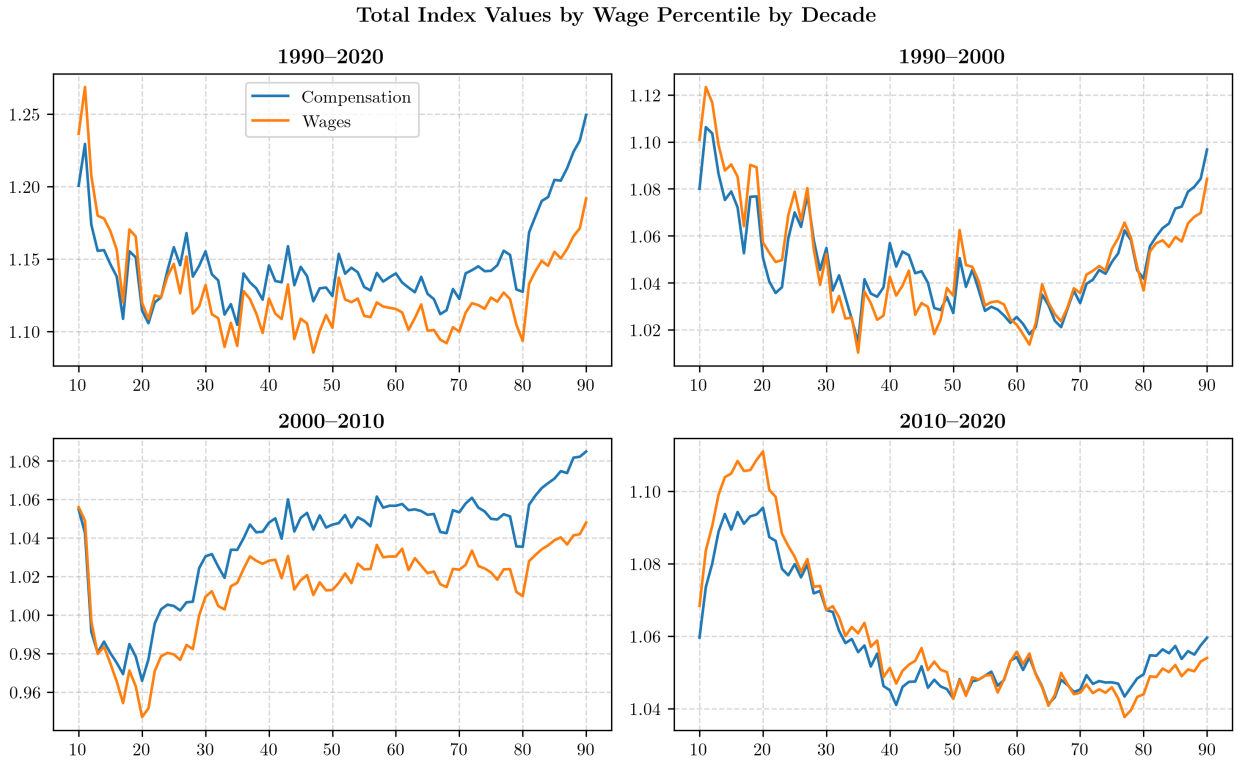


Figure 13: *Growth in Real Costs by Wage Percentile by Decade.* The above figure plots the total index values (I_t^{total}) for compensation and wage costs across the wage distribution for the 1990–2020 period and each intervening decade. We follow the approach of Pierce (2010) in our construction of these indexes. First, we sort all jobs by their average wage and use hours-adjusted sample weights to identify wage percentiles. Next, records are binned using a 20pp window centered on each wage percentile and then all indexes are computed for each cost component. In keeping with the chained structure of our index formulation, we re-bin the records for every quarterly change and then accumulate these changes over time.

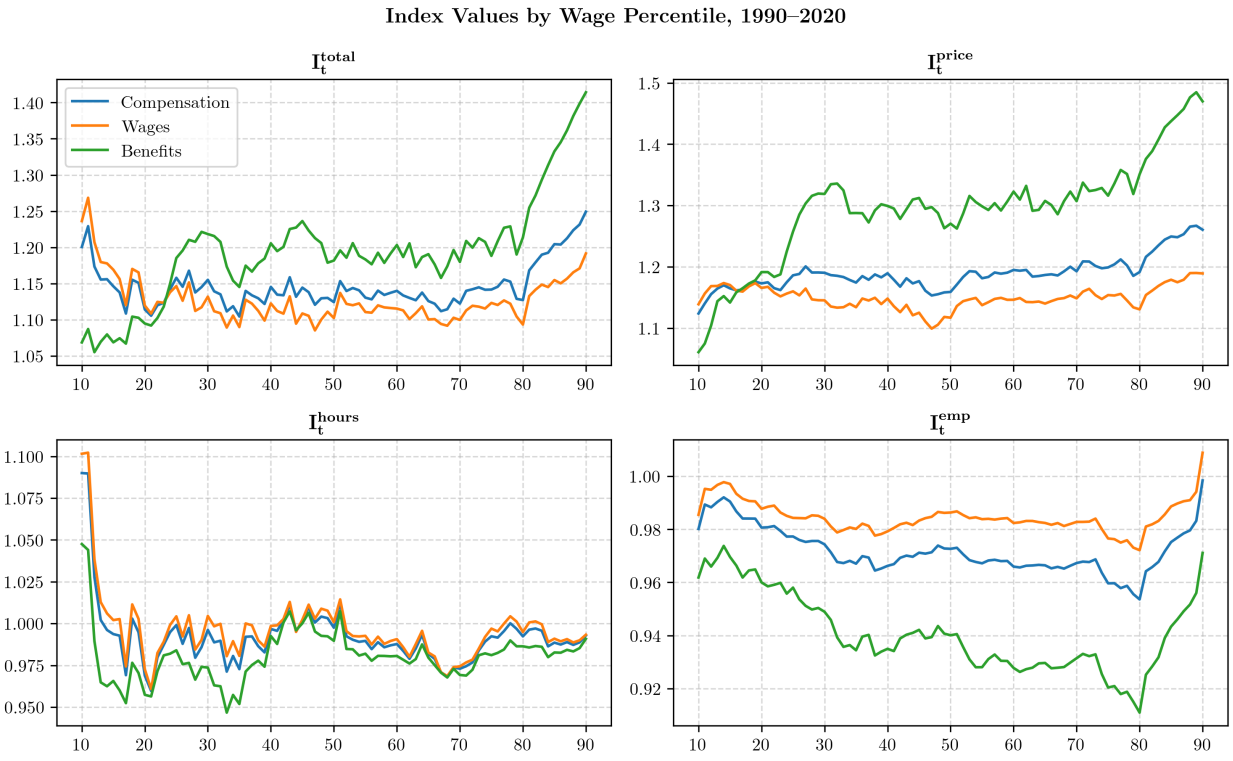


Figure 14: *Growth in Real Costs by Wage Percentile, 1990–2020.* The above figure plots the index values for compensation (including wages and benefits) costs across the wage distribution for the 1990–2020 period. The top left panel reports total changes, while the other three panels report the price, hours, and employment distribution decomposition. We follow the approach of Pierce (2010) in our construction of these indexes. First, we sort all jobs by their average wage and use hours-adjusted sample weights to identify wage percentiles. Next, records are binned using a 20pp window centered on each wage percentile and then all indexes are computed for each cost component. In keeping with the chained structure of our index formulation, we re-bin the records for every quarterly change and then accumulate these changes over time.

Online Appendix

C Weighting

C.1 Background

Survey programs must contend with at least two issues: (1) they are voluntary and (2) sampling methods may not represent the same population over time. The National Compensation Survey (NCS) addresses these concerns by adjusting its sample weights. This begins with the probability selection, which can be derived from the stratified sampling procedure and the sampling frame. Adjustment factors for non-response are then applied which redistributes weight to responding firms. In the NCS there are several levels at which this non-response is addressed: unit (establishment), quote (job), and item (cost within job). Finally, because weights are assigned when an establishment is sampled, adjustments are made at collection to reflect changes in the composition of employment.

The NCS calls this final step "benchmarking," though applied statisticians would call this weight "calibration" or (if they were willing to commit a minor faux pas) "post-stratification." The distinction between these latter two terms depends on the availability of a joint distribution with which to make adjustments (in which case the weights are being *post-stratified*) versus the use of multiple, separate marginal distributions (and thus *calibrated*). Bureau of Labor Statistics (BLS) data products fed by the NCS, for example the Employer Cost Index (ECI) and the Employer Costs for Employee Compensation (ECEC), handle this step differently. The ECI uses fixed employment weights using the Quarterly Census of Employment and Wages (QCEW; counts by industry) and Occupational Employment Statistics (OES; counts by occupation within industry). The ECEC, in contrast, makes a further adjustment using employment counts from the Current Employment Statistics (CES) so that employment weights are current across industry. The CES is used here due to issues about timing. QCEW and OES data come out at a sufficient lag so that these counts are not available early enough for the ECEC publication schedule.

To apply the benchmarking procedure, NCS records are first aggregated into *cells*. Exactly how these cells are constructed has changed over time, mostly due to changes in the way establishments (or workers) are classified by industry (or occupation). This issue underscores one challenge of this project: in order for pragmatic comparisons of estimates across time, weights should represent a comparable population to the greatest extent that is feasible. Depending on how far back one might want to make comparisons, the availability

of external data, along with data quality issues within the NCS itself, limits the scope and specificity of these adjustments. For more information about how weighting is performed in official NCS data products, see Bureau of Labor Statistics (2017). Next, we describe the weighting exercise for the present work.

C.2 Calculation

It will be useful to establish language to refer to the two broad eras of BLS compensation data. Without getting into too many details, the NCS was not always called "the NCS." Over time, as various collection efforts and data products related to compensation were combined, the current name was chosen. For ease of composition we refer to all of these data as the *NCS*, and distinguish between *historic* and *modern* vintages, where the modern era of NCS data begins in June 2004. Before this date industries are classified according to the Standard Industrial Classification (SIC) system, and after we observe North American Industry Classification System (NAICS) codes. For occupations we observe Occupational Classification System (OCS) codes in the early samples, and Standard Occupational System (SOC) codes in the later samples. These characteristics are key inputs to the grouping of jobs for final weight assignment. They inform how much preliminary weights are adjusted to reflect the changing distribution of employment.

How aggregated should industry, occupation, and establishment size groups be? It may be tempting to use the most specific level of classification (e.g. 6-digit NAICS for industry or 6-digit SOC for occupation) so that the implied distribution of employment matches the counts in our target distribution. However, there are no free lunches: if the cells defined by very specific classifications of jobs have a small number of records in them, variance issues become front and center. Aggregation limits this influence, and also alleviates some of the errors associated with converting between SIC and NAICS (or OCS and SOC). We use "PNAICS 8" groups when benchmarking weights by industry and "PSOC 5" for occupations.[11] Table A.1 and report how each of these aggregations is constructed. Size groups are defined by establishment employment in the following bins: 1–49 employees, 50–99 employees, 100–499 employees, and more than 500 employees.

The first external source we use come from the QCEW. This source provides employment counts by industry and size class, and is natural in that it is also the sampling frame of the NCS. These data are available the first quarter of every year, starting in 1990. Second, we use CPS employment counts for the PSOC 5 occupation groups which are available

[11] The "P" stands for "pseudo" and refers to custom aggregations of NAICS sectors or Major SOC groups.

through FRED starting in 1983. We note that the OES provides estimates of occupational employment by industry, but that the early period (including the 1990s) has many gaps: only a subset of industries is collected each year. In choosing these CPS estimates over the OES, we lose variation of changes to employment across occupations by industry, but gain a great degree stability over time. Noting the availability of data across these sources, we can calibrate weights beginning in 1990.

To construct benchmark adjustment factors for NCS jobs, we inflate or deflate preliminary weights so that domain sums equal control counts from target distributions. These factors are produced using the following raking procedure.

$$adj_{j,k,l} = \frac{ControlCount_j}{\sum_i w_{i,j,k,l}}(s_{j,l})(o_k), \tag{4}$$

i indexes jobs, j indexes PNAICS 8 industry group, k indexes PSOC 5 occupation group, and l indexes establishment size groups. $s_{j,l}$ is size-class l's employment proportion for industry group j, and o_k is occupation group k's employment proportion.

C.3 Assigning Occupation Groups to Historic Records

In samples before June 2004, the NCS program used OCS codes to group occupations. In the very earliest samples, detailed occupations are not always available. Instead, only "Major Occupational Groups" (MOGs) were recorded, which are the primary aggregates in the OCS. Unfortunately, while an OCS-to-SOC crosswalk exists for detailed codes, these MOGs do not cleanly map into the aggregates we need for weighting. For comparison, 2-digit SIC codes for industries are much more easily mapped into the 2-digit NAICS aggregates today. Thus, to apply the benchmarking method and control for shifts in occupations over time, we need to find a way to assign records those occupation groups (namely the PSOC 5). We use a probabilistic assignment model, where NCS covariates and observed OCS-to-SOC assignments are used to classify records with missing codes into the PSOC 5 occupation groups. Assuming a logit structure,

$$\Pr\left(OccGroup_i = k \middle| X_i, \theta_k\right) = \begin{cases} \exp(X_i'\theta_k)\left[1 + \sum_{k=1}^{K-1} \exp(X_i'\theta_k)\right]^{-1}, & k < K \\ \left[1 + \sum_{k=1}^{K-1} \exp(X_i'\theta_k)\right]^{-1}, & k = K, \end{cases} \tag{5}$$

where i indexes jobs, k indexes occupation groups, X are the covariates used to predict group assignment, and θ are parameters to be estimated with MLE. After merging all NCS data for which OCS codes exist (and therefore assignments into PSOC), we split the result into two sub-samples for training and validation. Letting a tilde denote the training sample for $i = 1, \ldots, \widetilde{N}$, we estimate

$$\widehat{\theta}_k := \underset{\theta_k}{\operatorname{argmax}} \sum_{i=1}^{N} \left\{ \widetilde{X}_i'\theta_k - \ln\left[1 + \sum_{k=1}^{K-1} \exp(\widetilde{X}_i'\theta_k)\right] \right\}, \quad k < K. \tag{6}$$

$\widetilde{X}$ includes a constant and dummies for MOG, PNAICS 8 industry group, establishment size, full/part-time status, time/incentive status, union status, state FIPS, hours-per-week quartiles, AHR quartile, and total benefit cost quartile, where AHR and benefit costs are inflation-adjusted using regional CPI figures. With those estimates, we then predict the probability of assignment into the PSOC 5 occupation groups using the validation sample. The results of this step show that this classification procedure correctly predicts group assignment 92% of the time.

D Validation with the ECI

In this section we compare our estimates with those of the Employment Cost Index (ECI), an official BLS data product utilizing the same underlying micro-data from the NCS. We begin by introducing the ECI and then discussing key differences between it and our exercise. From Bureau of Labor Statistics (2017), "The ECI is a measure of the change in the employer costs of labor, independent of the influence of employment shifts among occupations and industry categories." For a given domain (set of observations), an ECI estimate is constructed as follows.

$$W_{0,i} = \overline{Y}_{0,i} E_{0,i} \tag{7}$$

$$R_{t,i} = \frac{\overline{Y}_{t,i}}{\overline{Y}_{t-1,i}} \tag{8}$$

$$M_{t,i} = \prod_{s=1}^{t} R_{s,i} \tag{9}$$

$$I_t = \frac{\sum_i W_{0,i} M_{t,i}}{\sum_i W_{0,i}} \times 100, \tag{10}$$

where t identifies quarter and i names an ECI *basic cell*. Currently there are 539 Private sector basic cells generated from 59 industry groups (roughly 3-digit NAICS) by 9 occupational cells (an aggregation above 2-digit SOC). $W_{0,i}$ are the cost weights for cell i in the base period; $\overline{Y}_{t,i}$ is the weighted average cost in t for cell i; $E_{0,i}$ is the base period employment for cell i; $R_{t,i}$ is cell i's relative in t; $M_{t,i}$ is the multiplicatively accumulated relatives from the base-period to t for cell i; and I_t is the final ECI.

The above describes a *modified Laspeyres* index with fixed, base-period weights. The first difference we highlight is this weighting. The ECI's weights control the sample's implied distribution of employment by industry and establishment size (from the QCEW) and occupation (from the OES). This distribution is updated infrequently, roughly every 10 years, so estimates can be interpreted in pure-price, fixed-basket terms. In comparison, the weights used in our approach are chained, allowing the basket of labor services to gradually update over time. Chain indexes are better suited to measure changes to the "cost of living" (as opposed to pure-prices) because they allow substitution to occur. The basket structure also differs in the way observations are aggregated for calibration to facilitate the combining of samples across different coding regimes (SIC to NAICS for industry; OCS to SOC for occupation).[12]

We also note the *order* of aggregation is different between our estimates and the ECI. The basic cells of the ECI prescribe an intermediate aggregation stage where growth relatives are first calculated for a cell before being combined across cells to give a final result. This paper's indexes calculates totals across cells (where cell weights are loaded into the calibration process) before constructing relatives. Our approach is closer to a textbook treatment of a Laspeyres price index and, importantly, can be combined with appropriately defined quantity

[12]For additional details see Appendix C.

indexes to enable for the decomposition of total changes by prices, hours, and the employment distribution separately.

Despite the above it is still useful to compare our estimates with those of the ECI, where possible. Though some differences affect which domains are available, there is still limited overlap in series. Further, because the ECI is concerned with pure price changes, we present ECI series alongside our price sub-indexes (I_t^{price}). **??** reports our estimates (solid line) and the ECI estimates (dotted line) for compensation (blue), wages (orange), and benefits (green). Constant dollar ECI estimates for Private sector workers are available beginning in the first quarter of 2001. The series are broadly similar, perhaps most notably different for benefits. Our estimates suggest a higher growth rate to the tune of 2–3pp from 2001 to 2023. We also plot the comparison for occupation and industry domains in **??**. The estimates diverge more than with the aggregate series, but the broad story generally holds. We note that the largest differences appear for the Trade sectors: our estimates for compensation and wages are higher for Wholesale Trade but lower for Retail Trade.

E Tables

PNAICS 8	**2-digit NAICS**	**2-digit SIC**
Mining	21	10–14
Construction	23	15–17
Manufacturing	31–33	20–39
Transportation, Communications, and Utilities	22, 48–49, 51	40–49
Wholesale Trade	42	50–51
Retail Trade	44–45	52–59
Finance, Insurance, and Real Estate	52–53	60–67
Services	54–56, 61–62, 71–72, 81	70–89
PSOC 5	**2-digit SOC**	
Management, Professional, and Related	11–29	
Sales and Office	41–43	
Service	31–39	
Natural Resources, Construction, and Maintenance	45–49	
Production, Transportation, and Material Moving	51–53	
Size Class		
0–49 Employees		
50–99 Employees		
100–499 Employees		
500+ Employees		

Table A.1: *Weight Calibration Groups.* This table reports the industry, occupation, and establishment size aggregates used in weight calibration.

	1990–2000			2000–2010			2010–2020		
	Comp.	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*
Industry Group									
Mining	1.6	-1.1	8.4	20.1	9.5	45.6	7.8	10.2	3.5
Construction	2.5	0.5	8.1	8.2	5.8	14.2	6.3	6.3	6.3
Manufacturing	6.9	6.0	9.0	6.8	2.4	16.8	7.1	6.6	8.2
Transp., Comms., and Utilities	5.8	3.1	12.4	8.0	3.3	19.4	9.0	7.7	11.9
Wholesale Trade	8.3	6.4	13.3	7.0	2.9	17.8	9.2	8.5	11.0
Retail Trade	7.3	7.0	7.8	5.3	0.9	21.3	5.6	5.9	4.7
Finance, Insurance, and Real Estate	11.6	9.7	16.7	9.4	4.8	21.3	8.4	8.4	8.6
Services	7.7	6.6	11.2	7.6	4.5	17.0	4.6	4.7	4.6
Size Group									
0-49 Employees	7.7	6.6	11.1	4.3	2.0	11.5	5.1	5.4	4.3
50-99 Employees	8.1	7.6	9.4	8.4	5.1	17.6	6.6	6.0	8.3
100-499 Employees	7.6	6.6	10.4	9.1	4.9	20.3	7.7	7.4	8.6
500+ Employees	5.5	3.2	10.6	10.9	4.6	25.9	5.9	5.3	7.3
Occupation Group									
Management, Professional, and Related	8.3	6.6	13.1	6.9	3.4	16.6	5.2	4.9	6.0
Sales and Office	9.2	8.8	10.4	9.3	5.1	21.0	8.8	8.8	8.7
Service	3.1	1.3	7.8	7.5	2.3	21.2	8.6	9.3	7.0
Nat. Res., Constr., and Maintenance	5.4	4.6	7.3	9.5	4.8	21.6	4.7	3.3	8.0
Production, Transp., and Mat. Moving	3.8	2.2	7.4	5.8	2.1	15.4	8.0	8.7	6.6

Table A.2: *Group Cost Growth from Prices.* This table reports the real growth rates in costs from price changes (I_t^{price}) broken down by industry, establishment size, and occupation. For example, the real cost of compensating Sales and Office workers increased by 9.3% in the 1990s. For these workers, wage costs grew by 8.8% and benefit costs increased by 10.4%. These group-level estimates will roughly aggregate to the overall price index estimates in Table 1 using weights for the relative shares of each group, but may not be exact due to rounding and the chained aspect of the index (that is, weights are updated each period).

	1990–2000			2000–2010			2010–2020		
	Comp.	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*
Industry Group									
Mining	-3.5	-3.7	-3.0	-1.1	-0.9	-1.4	-0.5	-0.5	-0.5
Construction	-0.5	-0.2	-1.4	-0.4	-0.6	-0.0	-0.3	-0.3	-0.3
Manufacturing	0.1	0.2	-0.2	-1.2	-1.1	-1.3	-0.5	-0.6	-0.4
Transp., Comms., and Utilities	-2.2	-1.9	-2.9	-0.4	-0.2	-0.8	0.0	-0.1	0.2
Wholesale Trade	-0.6	-0.7	-0.5	-1.7	-1.7	-1.6	0.3	0.3	0.1
Retail Trade	-0.0	0.5	-1.7	0.2	0.4	-0.2	0.1	0.2	-0.2
Finance, Insurance, and Real Estate	-0.7	-0.8	-0.6	2.3	3.0	0.6	-0.7	-0.7	-0.7
Services	1.0	1.1	0.6	-0.5	-0.4	-0.6	-0.6	-0.6	-0.8
Size Group									
0-49 Employees	1.1	1.3	0.6	0.6	0.8	0.1	-0.6	-0.5	-0.6
50-99 Employees	-2.1	-2.0	-2.4	-1.1	-0.8	-1.8	-0.2	-0.1	-0.3
100-499 Employees	-0.5	-0.3	-1.1	-1.1	-1.1	-1.2	-0.4	-0.4	-0.5
500+ Employees	-0.2	-0.1	-0.5	-0.7	-0.8	-0.6	-0.3	-0.3	-0.3
Occupation Group									
Management, Professional, and Related	0.1	0.1	-0.1	0.2	0.4	-0.3	-0.4	-0.4	-0.5
Sales and Office	-0.6	-0.4	-1.2	-0.9	-0.8	-0.9	-0.3	-0.3	-0.2
Service	-0.9	-0.2	-2.6	-1.7	-1.6	-2.0	-0.1	-0.0	-0.3
Nat. Res., Constr., and Maintenance	-1.1	-1.1	-1.1	-0.9	-0.8	-1.0	-1.1	-1.0	-1.2
Production, Transp., and Mat. Moving	2.0	2.4	1.0	-0.7	-0.7	-0.7	-0.0	-0.0	0.0

Table A.3: *Group Cost Growth from Hours.* This table reports the real growth rates in costs from hours changes (I_t^{hours}) broken down by industry, establishment size, and occupation. For example, the real cost of compensating Sales and Office workers decreased by 0.6% in the 1990s due to reduced hours. For these workers, wage costs fell by 0.2% and benefit costs decreased by 1.2%. This suggests that hours fell by more in jobs that are relatively benefit heavy.

	1990–2000			2000–2010			2010–2020		
	Comp.	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*	*Comp.*	*Wages*	*Benefits*
Industry Group									
Mining	-0.3	0.2	-1.3	2.5	2.3	2.7	0.6	0.9	0.3
Construction	2.4	2.3	2.5	1.2	1.3	1.0	5.7	5.4	6.7
Manufacturing	1.0	1.6	-0.1	0.0	0.7	-1.4	4.1	3.9	4.6
Transp., Comms., and Utilities	-0.8	-0.2	-2.1	-0.0	0.2	-0.5	5.1	5.0	5.3
Wholesale Trade	1.2	1.5	0.5	1.9	2.1	1.4	3.1	2.9	3.5
Retail Trade	1.4	1.5	1.1	3.7	3.8	3.4	5.2	5.1	5.4
Finance, Insurance, and Real Estate	2.7	3.1	1.8	2.0	2.2	1.5	4.0	3.7	4.7
Services	1.6	1.8	1.2	1.9	2.1	1.4	3.4	3.3	3.5
Size Group									
0-49 Employees	1.3	1.4	1.1	1.2	1.4	0.7	2.9	2.8	2.9
50-99 Employees	1.0	1.2	0.3	0.1	0.5	-0.9	2.9	3.0	2.9
100-499 Employees	0.0	0.5	-1.2	-0.3	0.4	-1.9	3.4	3.6	3.1
500+ Employees	-0.5	-0.1	-1.2	-1.7	-0.9	-3.3	3.3	3.7	2.7
Occupation Group									
Management, Professional, and Related	-2.6	-2.4	-3.1	-2.8	-2.3	-3.9	0.2	0.1	0.3
Sales and Office	-2.4	-1.9	-3.5	-3.8	-3.1	-5.6	0.5	0.3	0.9
Service	-4.9	-3.7	-7.8	-7.0	-5.4	-10.6	0.5	0.3	1.1
Nat. Res., Constr., and Maintenance	-1.8	-1.1	-3.3	-2.3	-1.4	-4.5	0.7	0.5	1.1
Production, Transp., and Mat. Moving	-3.6	-2.8	-5.5	-5.5	-4.4	-8.2	0.7	0.6	1.0

Table A.4: *Group Cost Growth from Employment Shifts.* This table reports the real growth rates in costs from changes to the distribution of employment (I_t^{emp}) broken down by industry, establishment size, and occupation. For example, the real cost of compensating Sales and Office workers decreased by 2.4% in the 1990s due to shifts of employment into different industries and establishment sizes. For these workers, wage costs fell by 1.9% and benefit costs decreased by 3.5%. This suggests that employment of Sales and Office workers has shifted away from industries and establishment sizes that are relatively benefit heavy.

F Figures