Composition-Adjusted Wage Growth:

A Robust Measure from Microdata

Bo E. Honoré and Luojia Hu

July 2025

WP 2025-14

https://doi.org/10.21033/wp-2025-14

FEDERAL RESERVE BANK *of* CHICAGO

Composition-Adjusted Wage Growth: A Robust Measure from Microdata*

Bo E. Honoré† Luojia Hu‡

July 2025

Abstract

Wage growth is a key indicator of labor market conditions, but common measures often conflate individual wage changes with shifts in workforce composition. This paper develops a composition-adjusted measure of wage growth using nonparametric decomposition and program evaluation methods. The adjusted measure tracks unadjusted growth in stable periods but diverges during disruptions: during the Covid-19 pandemic, wage growth falls from 12% to 6% after adjustment. The method accommodates rich covariates, is robust to data quality issues such as rounding, heaping and top-coding, and enables distributional and subgroup analysis using micro data, offering more accurate views of underlying wage dynamics.

Keywords: Wage Growth, Selection, Decomposition, Robust Measures.

JEL Codes: J31, C21, C18.

---

* **Acknowledgments**: The research was supported by the Gregory C. Chow Econometric Research Program at Princeton University, by the National Science Foundation (Grant Number SES-1530741) and by the Danish National Research Foundation grant number DNRF186 to the Aarhus Center for Econometrics (ACE). We thank Kristin Butcher, Daniel Sullivan, participants at the Princeton Industrial Relations Section's Centenary Symposium and the Federal Reserve Bank of Chicago Brown Bag Seminar for helpful suggestions and comments. The opinions expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

***Conflict of Interest Statement**: The authors have no conflict of interest to declare.

†Princeton University, USA. Email: honore@princeton.edu.

‡Federal Reserve Bank of Chicago, USA. Email: lhu@frbchi.org.

1 Introduction

Discussions of macroeconomic conditions frequently rely on summary indicators that capture specific dimensions of economic activity. In the context of the labor market, commonly cited measures of wage growth include the Federal Reserve Bank of Atlanta's Wage Growth Tracker, the percentage change in the Bureau of Labor Statistics' (BLS) Average Hourly Earnings (AHE), and the percentage change in the BLS Employment Cost Index (ECI). These indices are widely used in both policy and academic settings to assess trends in earnings and labor market tightness. Figure 1 illustrates the trends in these indicators over time. While the three measures followed similar patterns through 2020, they have diverged significantly since then. This divergence stems from differences in survey sources and target populations. Notably, the Wage Growth Tracker and the ECI adjust for changes in workforce composition, whereas the AHE does not. Each measure relies on distinct data sources: the Atlanta Fed's Wage Growth Tracker uses the Current Population Survey (CPS) and focuses on workers with a job in both of two months one year apart; the AHE is based on data from the Current Employment Statistics (CES) survey, which samples employers; and the ECI uses the National Compensation Survey (NCS), tracking compensation for a fixed set of industry-occupation groups.

In this paper, we use nonparametric decomposition and program evaluation methods to construct an alternative composition-adjusted measure of wage growth. Like the Atlanta Fed's Wage Growth Tracker, we use data from the CPS, but unlike the Atlanta Fed's Wage Growth Tracker, we focus on all workers at a point in time, not only the subset of workers who are employed in both of two periods twelve months apart. Moreover, our analysis does not exclude top-coded wages, and it aims to be robust to other common data quality issues such as rounding, heaping, and to outliers. Like the ECI, we hold some characteristics of the jobs fixed. However, unlike the ECI, we also control for worker characteristics.

To motivate and illustrate our approach, the average wage rate in our sample in July 2019 was \$25.39 while it was \$27.82 in July 2020. This translates to a 9.6 percent increase. Although this increase could be a result of increases in wages at the individual level, it could also reflect a change in the composition of the wage earners in the direction of more higherpaid workers. This makes it of interest to decompose the overall wage change into the part that corresponds to the a shift in the distribution of characteristics (holding the average wage for a given set of characteristics fixed at the initial level) and the part that corresponds to a shift in the distribution of wages given the characteristics (holding the distribution of characteristics fixed). This is illustrated in Table 1, which reports the distribution of educational attainment among workers in each of the two periods. The table shows that the workforce has shifted towards higher educated workers and that in each period, more educated workers have higher wages than less educated workers. As a result, average wages would have increased even if the average wage rate remained the same in each of the education categories. Using the numbers from Table 1, the increase that one would have seen without an increase in the average wage within an education group, would be 3.4 percent. Alternatively, one could calculate the wage rate that would have materialized if the distribution of education had been unchanged between 2019 and 2020. This number is 6.1 percent. Both of these numbers (3.4 and 6.1 percent) are economically quite different from the 9.6 percent raw increase in the average wage, and one can argue that the number that holds the distribution of education constant over time is a better measure of aggregate wage growth.

Table 1: Distribution of Wages Across Education in 2019 and 2020

	Less than HS	High School	Some College	College	Advanced	All
2019	\$14.09 (7.08%)	\$18.56 (25.29%)	\$21.98) (29.60%)	\$31.31 (24.8%)	\$40.93 (13.22%)	\$25.37 (100%)
2020	\$14.13 (6.02%)	\$19.69 (23.33%)	\$23.75 (27.44%)	\$33.72 (27.40%)	\$42.04 (15.70%)	\$27.82 (100%)

Note: For 2019 and 2020, the table reports the average hourly wage by educational category and (in parenthesis) the corresponding share of the education category. The data is from IPUMS CPS, and the table covers individuals aged 16-70.

The calculations presented in the example above are straightforward and well understood when one wants to hold fixed the distribution of a few categorical variables (such as educational attainment). However, in practice, one would want to control for a large set of characteristics of the workforce. These include demographic characteristics like age and sex as well as job characteristics such as industry and occupation. In a linear regression context, one could use textbook Kitagawa–Blinder–Oaxaca techniques to account for composition effects. However, issues related to rounding, heaping, top-coding and sensitivity to outliers make this approach unattractive in our context. We therefore turn to techniques from the treatment effect estimation literature such as propensity score matching, reweighting, and distributional regression. In order to allow for a significant number of covariates, parts of the implementation of the methods make use of penalized logit estimation.

Throughout the paper, we focus on annual wage growth and fix the distribution of the characteristics at the distribution in our sample at the end of the twelve month span. In other words, we measure the wage growth over the past twelve months for a set of workers like the ones who are currently working.

This paper contributes to the applied econometrics literature on program evaluation and decomposition methods. See, for example, Imbens and Wooldridge (2009a) and Fortin, Lemieux, and Firpo (2010) for classic reviews of these areas. It also engages with a more policy-oriented literature concerned with the measurement of wage growth; see, for example, Aaronson and Sullivan (2001) and Cole, Hu, and Schulhofer-Wohl (2017).

Figure 2 previews the findings of our paper. This figure shows the unadjusted annual percentage change in our preferred summary statistic (the 15-85% trimmed mean) for wages in blue, while the red curve displays a version of the growth in the same measure that has been adjusted for the change in the composition in the labor force from one year to the next. (See Section 4.) During "normal" times, our composition-adjusted wage growth measure tracks the unadjusted measure of wage growth, but the two differ dramatically in periods in which the composition of the workforce changes significantly. This is most pronounced during and immediately after the Covid pandemic. During this period, peak estimated wage growth falls from 12% to 6% after adjustment, while a trough of -2% becomes nearly $+2\%$.

We describe the data in Section 2. In Section 3, we discuss some of the challenges associated with the data. This includes rounding, heaping, and top-coding. Section 4 presents our methodology and the main results. In Section 5, we compare the composition-adjusted wage growth measure across gender, across educational groups, and across different parts of the wage distribution. Section 6 discusses additional data related issues such as the use of sampling weights. Section 7 concludes.

2 Data

For the analysis in this paper, we use the Current Population Survey (CPS) Basic Monthly micro data from January 2006 to December 2024. The CPS is widely used for calculating wage growth, such as in the Federal Reserve Bank of Atlanta's wage-growth tracker. The CPS is attractive due to its representative nature and detailed information on individual demographics and job characteristics, including industry and occupation. In addition, new waves of data are released monthly, which allows for almost real-time calculation of the wage growth. The data is publicly available, and institutions like IPUMS provide user-friendly versions (see Flood, King, Rodgers, Ruggles, Warren, Backman, Chen, Cooper, Richards, Schouweiler, and Westberry (2024)).

The CPS has a panel design where households are interviewed for four consecutive months, then not interviewed for eight months, and finally interviewed for another four consecutive months. Although many variables are collected each interview month, earnings-related questions are only asked in the outgoing rotation groups (the fourth and eighth interview months), so our sample is restricted to those in these "outgoing rotation" groups. We further restrict our sample to individuals aged 16 to 70 (inclusive) and exclude military personnel. For more details, see the Data Appendix.

3 Challenges Resulting from The Nature of the Data

Beyond economic considerations (see, for example, Cole, Hu, and Schulhofer-Wohl (2017)), there are also important statistical reasons to favor certain measures of wage growth over others. The composition-adjusted wage growth measure presented below uses a trimmed mean of earnings. This section explains why we adopt this approach.

Extremely high earnings values can sometimes make individuals identifiable within a dataset. To protect anonymity, many publicly available datasets—such as the CPS—apply top-coding. This means any earnings above a certain threshold are reported at that threshold instead of their true value. Sometimes, a separate flag indicates that a value has been top-coded. As a result, average earnings calculated directly from these datasets may be misleading. This issue becomes more pronounced when analyzing changes over time: unless the top-coding threshold and the reported top-coded value increase in line with overall wage growth, part of the observed change in average wages may simply reflect the effects of top-coding. Figure 3 illustrates this concern. The blue line shows wage growth using the average of all reported wages, which indicates a sharp increase starting in spring 2023. However, this spike does not appear in the $15-85\%$ trimmed mean. The discrepancy is due to a change in how the Bureau of Labor Statistics handles top-coded wage data — highlighting the importance of either dealing explicitly with top-coding or reporting measures that are robust to top-coding.

Researchers have also raised concerns about the reliability of publicly available survey data on wages and earnings when it comes to the lower end of the distribution. See for example Rinz and Voorheis (2018). A key issue is misreporting—some individuals may report implausibly low hourly wages, such as five cents or one dollar. Additionally, changes in minimum wage laws can make changes in low-end wages less indicative of underlying economic conditions than movements in the middle of the distribution.

Given these challenges, focusing on median wages might seem preferable to using averages. However, medians have limitations of their own. Reported wages often cluster at round numbers due to actual wages heaping on round numbers and due to respondents' tendency to round. This introduces a form of discreteness that can obscure meaningful variation. For instance, if the median hourly wage is reported as $20 and many individuals round their wage to that number, significant shifts in the true distribution of wages may occur without affecting the median.

To address these issues, we use trimmed means as a more robust summary measure.[1] A trimmed mean calculates the average of a variable after removing a specified percentage of

[1] We are not the first to apply trimmed means to improve the robustness of economic indicators. For instance, the Federal Reserve Bank of Dallas uses a trimmed mean to calculate a core measure of price inflation. See https://www.dallasfed.org/research/pce.

the lowest and highest values. For percentiles p_1 and p_2, this means discarding the bottom p_1% and the top $(100 - p_2)$% of observations before computing the mean. Since fewer than 10% of earnings are top-coded each month, we use a 15–85% trimmed mean in our analysis. For comparison, the Atlanta Fed's Wage Growth Tracker simply excludes top-coded earnings values.

4 Composition Effects

One challenge in using observed wages to measure wage inflation is that the composition of the workforce can change over time. For instance, during an economic expansion, individuals who are only marginally attached to the labor force may enter employment. If these new entrants tend to earn lower wages, aggregate wage measures could decline – even if wages are rising for those already employed.

To address this, some measures of wage growth focus on sub-sample of individuals who are employed in both periods. This approach, which is, for example, used by the Atlanta Fed's Wage Growth Tracker, helps isolate true wage changes by holding workforce composition constant.

This method is not without limitations. Comparing the same individuals across two months that are one year apart inherently involves comparing groups that differ in age by one year, which may introduce bias if wages are age-dependent. More substantively, the subset of individuals who remain employed in both periods may not be representative of the broader labor market, thereby limiting the external validity of the findings (see, for example, Cole, Hu, and Schulhofer-Wohl (2017)). Another important limitation concerns the ability to reliably link individuals across time. Although the CPS is structured as a panel, the publicly available versions lack a unique personal identifier. Researchers must therefore construct identifiers using a combination of available variables — a process that is prone to mismatches and potential measurement error. Finally, restricting attention to individuals observed as employed in both periods necessarily confines the analysis to wage changes over a 12-month horizon.

4.1 Selection into the Sample of Wage Earners

The problem alluded to above is that individuals are not randomly selected into the sample of wage earners. In a given period, some people are more likely to work than others and who they are changes over time.

There are two types of selection that are relevant for studying wage growth using survey data: selection into the sample and selection into working. For both, it is potentially important to consider selection on observables as well as selection on unobservables. With selection on observables, the selection is allowed to depend on an individual's observed characteristics, but not on unobserved characteristics that are related to wages; see, for example, Heckman, Ichimura, and Todd (1997, 1998). Handling selection on unobservables is much more challenging than selection on observables, and doing so often requires one to make exclusion assumptions that cannot always be tested. For a discussion of estimation of sample selection models without exclusion restrictions, see, for example, Escanciano, Jacho-Chávez, and Lewbel (2016) and Honoré and Hu (2020). For this reason, this paper focuses on selection on observables.

In order to adjust wage growth for composition effects, one must compare wages for workers in one period to wages for workers with the same characteristics in a previous period. By comparing earnings for this "matched" sample in period $t-1$ to the earnings in period t, one obtains an apples-to-apples comparison of earnings in the two periods. To do this, one must ask what wage distribution workers from this period would have had last period. Denoting wages and characteristics as y and x, respectively, this makes the object of interest

$$\int f_{t-1}\left(y|x\right) f_t\left(x\right) dx \tag{1}$$

where $f_{t-1}\left(y|x\right)$ is the distribution of wages in period $t-1$ for an individual with characteristics x and $f_t\left(x\right)$ is the distribution of the characteristics in period t. One can then calculate summary measures for this "artificial" wage distribution and compare them to the same measures for the actual distribution in period t. As discussed above, this paper focuses on the trimmed mean as the most interesting measure.

There are a number of ways to approach the construction of the distribution in (1), and our final measure will be a combination of some of these.

4.2 Kitagawa–Blinder–Oaxaca Decomposition

It is straightforward to construct the mean of the distribution in (1) in a regression setting. Specifically, in a classical linear regression model, the coefficients contain all the information that one needs to estimate the expected value of the dependent variable for an individual with specific characteristics (such as age, gender, occupation, etc.). One can therefore decompose the difference in means between two groups into the difference that comes from different regression coefficients and the difference that comes from different means of the explanatory variables. This is done by first estimating the model using one sample, and then combining the estimated coefficients with the explanatory variables from the second sample. This is known as a Kitagawa–Blinder–Oaxaca decomposition. See Oaxaca and Sierminska (2023) for a discussion of the similarities and differences in the measures proposed by Kitagawa (1955), Blinder (1973), and Oaxaca (1973).

Two complications prevent us from doing a simple Kitagawa–Blinder–Oaxaca in our context. The first is that the censoring induced by top-coding (and the potential problems with the left tail of the wage distribution) can make it inappropriate to model wages by a linear regression model (Tobin (1958)).

The second problem is that here the trimmed mean is the object of interest and linear regression focuses on the conditional mean. In other words, one cannot decompose the difference in the trimmed mean between two groups by simply estimating a regression model.

4.3 Distributional Regression

The cumulative distribution function for the distribution in (1) is $\int F_{t-1}(y|x)\, f_t(x)\, dx$ where $F_{t-1}(y|x)$ is the cumulative distribution function for wages in period $t-1$ for individuals with characteristics x. A nonparametric version of the Kitagawa–Blinder–Oaxaca decomposition of wages would therefore first estimate the cumulative distribution of wages conditional on the characteristics in period $t-1$ and then average this over the distribution of characteristics in period t. In principle, one could estimate $F_{t-1}(y|x) = P_{t-1}(Y \leq y|x)$ nonparametrically using data from period $t-1$. However, since the number of characteristics will be relatively large, this will be infeasible. Using data from period $t-1$, we therefore use a logit model to estimate the probability of the event $\{Y \leq y\}$ as a function of the characteristics. Doing this for a range of values of y will trace out an estimate of the conditional wage distribution. This approach is not new, and it is often referred to as distributional regression (Foresi and Peracchi (1995)). The idea of using distributional regression to estimate counterfactuals is also not new (Chernozhukov, Fernández-Val, and Melly (2013)).

Averaging the conditional wage distribution in period $t-1$ over the distribution of the characteristics in period t, gives an estimate of the cumulative distribution function for the distribution in (1), and this distribution can then be used to calculate summary measures for the "artificial" wage distribution in (1).

4.4 Matching

An alternative approach to measuring wage growth is to consider being observed in year t as opposed to year $t-1$ as a treatment. If we are interested in the wage growth for workers in year t, then the issue becomes one of how to measure the effect of a treatment on the treated. This problem has a long history in economics and statistics (see, for example, Imbens and Wooldridge (2009b)), and one approach is matching. The basic idea is to create an artificial sample of wages as follows: for each worker in time period t, find a worker with the same observed characteristics in time period $t-1$, and then include that worker's wage in the artificial sample. This will create a sample of wages from time $t-1$ for workers with the same characteristics as the workers in time t. Finding a worker from the sample in $t-1$ that has the same (or even similar) observed characteristics as a given worker in time t, can be difficult or impossible if one considers many observed characteristics. A central result in this literature is that rather than matching on the full vector of characteristics, one can match on the propensity score (Rosenbaum and Rubin (1983)), where "propensity score" refers to the probability of being treated as a function of an individual's characteristics.

To implement propensity score matching in the context of wage growth, we start with all workers in periods t and $t-1$, and estimate a logit model for whether the worker is from period t. This gives our estimated propensity score. We then use the following procedure to construct a counterfactual distribution for wages for period $t-1$ for a population that has (approximately) the same distribution of the propensity score as in period t: For each individual, i, in period t, we find the 200 observations[2] from period $t-1$ that have propensity scores closest to the propensity score for i. In the counterfactual distribution, these are given weights that depend on how close the propensity score is to that of observation i.[3] We then use the resulting distribution to calculate trimmed mean (or some other measure of location) and compare it to the one calculated from the sample of workers in time period t.

4.5 Reweighting

Rather than creating a counterfactual sample of wages in period $t-1$ by matching observations based on their characteristics or their propensity scores as above, one can reweight the observations in time period $t-1$ in such a way that the reweighted sample is similar to the sample in time period t in terms of observed characteristics. It is again useful to consider the probability that an individual with characteristics x belongs to period t (as opposed to $t-1$). This is what we referred to as the propensity score above. Here we denote it as $p_t(x)$, and $p_{t-1}(x)$ will denote the probability that an individual with characteristics x belongs to period $t-1$. Using simple manipulations of conditional probabilities, DiNardo, Fortin, and Lemieux (1996) show that the weights that will make the reweighted sample of observations in time $t-1$ similar to the sample in time period t in terms of observed characteristics are

$$w(x) = \frac{p_t(x)}{p_{t-1}(x)} \frac{p_{t-1}}{p_t}$$

where p_{t-1} and p_t are the proportions of observations in period $t-1$ and t, respectively.

[2]There are two exceptions to this. Since there can be ties in the estimated propensity scores, we can in principle use more than 200 "neighbors". Moreover, we exclude propensity scores that differ by more than 0.1. If this results in no "matches" then we use the closest. (There can again be more than one.)

[3]The weights add up to 1. The weight is quadratic in the difference in the propensity scores, scaled so that an exact match (i.e., an observation with exactly the same propensity score) results in a weight that is n times that of the lowest weight (the weight of the poorest match), where n is the number of matches used.

4.6 Implementation and Results

In this paper, we pursue the distributional regression, matching and the reweighting approaches to estimate composition-adjusted annual wage growth for each month.[4] For each of these, we calculate the three month moving average, and present the average of those moving averages as the preferred measure. Since the data is monthly, and we are interested in annual growth, the lagged period will be $t - 12$ rather than $t - 1$.

For the first measure, which is based on distributional regression, we use the data from period $t - 12$ to estimate a logit model for the probability that the wage rate is less than or equal to a number of values, y. For y, we use all of the percentiles of the wage rate in period $t - 12$, and the explanatory variables are age and age-squared as well as indicators for month in sample, sex, level of education, presence of children in the household, presence of children under 5, marital status, US-citizenship, race/ethnicity group, industry and occupation. For education, we allow for five categories: less than high school, high school graduate or equivalent, some college, college degree, and advanced degree. Our mutually exclusive race/ethnicity groups are White, Black, Hispanic, Asian and other. We aggregate industry and occupation into 14 industry categories and 26 occupation categories. This gives a total of 68 explanatory variables, some of which are linear combinations of the others. Estimating the logit model with maximum likelihood can lead to imprecision when there are many explanatory variables or when the set of explanatory variables includes "thin" categorical variables. We therefore estimate the logit model using maximum penalized likelihood with a ridge penalty and the penalty parameter chosen by cross validation.[5]

For each percentile and for each observation in period t , we then calculate the predicted

[4] The methods used here could be used to construct wage growth over arbitrary time spans. We focus on twelve month wage growth in order to avoid seasonality issues.

[5] For given penalty parameter, λ, the objective function is

$$\min_{\beta_1,\beta_2} \left\{ -\frac{1}{N} \sum_{i=1}^{N} [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] + \lambda \|\beta_2\|^2 \right\} \tag{2}$$

where the p_i's are the probabilities implied by the logit model. β_1 refers to the constant and the coefficients on age, age-square and dummy variables for the month in interview. β_2 refers to the coefficients on the remaining covariates. Only β_2 is penalized. The penalty parameter λ is chosen by 10-fold cross validation. All estimation was done in Stata Version 18.

probability (using the model estimated in period $t-12$) and average these across the observations. This gives an estimate of $\int F_{t-12}(y|x) f_t(x)\, dx$. The differences in these estimated probabilities between two adjacent percentiles will be the probability masses in a discrete approximation to the distribution in (1). For the corresponding points of support, we use the average wage rate (in period $t-12$) for the observations between the percentiles.

For the two other measures, the first step is to estimate the propensity score. To do this, we construct a sample of observations from periods t and $t-12$, and we then estimate the probability that an observation in this sample belongs to period t as a function of the observed characteristics. The observed characteristics that we use are the same as above, and we again used a logit model with Ridge penalization to estimate the parameters in the model for the propensity score.

Figure 4 displays the time series of the three adjusted wage growth measures. The figure confirms the idea that the three are supposed to be different estimates of the same counterfactual. In order to reduce estimation noise, we focus on the average of the three.

Our proposed composition-adjusted wage rate growth measure is depicted in Figure 5, which also shows the growth in the unadjusted trimmed mean. Except that it covers a longer time period, including the Great Recession, it is the same figure as Figure 2 in the introduction.

Figure 5 suggests that in "normal" times when there are no large movements in employment, the adjustment makes little difference for the measured wage growth. This is to be expected since such periods are unlikely to see great changes in the composition of workers. In contrast, the Great Recession and the Covid pandemic period were associated with large changes in employment. In the Covid period, in particular, the change in employment rates differed dramatically across groups, which would lead to shifts in the composition of workers. For example, Honoré and Hu (2023) found that minorities and individuals with lower educational attainment experienced larger decreases in employment at the onset of the pandemic. To the extent that these were also individuals with lower wages, this would lead to higher overall wage rates even if no single individual had a change in their wage. This is very visible in Figure 5. The annual increase in the 15%-85% trimmed mean of the observed wage rates was as high as 12%. After adjusting for the composition of wage earners, this was reduced by a factor of two. The economic recovery at the later stage of the pandemic illustrates the opposite effect. As the lower wage earners returned to work, the aggregate measures of wage growth fell (and even became negative, reaching a value of -2%), but it is clear from Figure 5 that this is partly due to a change in the composition of workers. After adjusting for this change in composition, the fall in wage growth becomes much less dramatic (and the adjusted wage growth is always positive).

Although less dramatic, the Great Recession (December 2007 to June 2009) illustrates the same points. Unemployment increased sharply; if the affected workers have lower wages than the ones that remained employed, then this will increase aggregate measures of wage growth even if the wage of no single continuously employed worker increased.

5 Heterogeneity in Wage Growth

One advantage of using individual level micro data is that this allows one to estimate wage growth for different groups. In Section 5.1 below, we do this by comparing wage growth for men and women and by comparing wage growth for different educational groups. Measuring wage growth by trimmed means also makes it possible to investigate how the wage growth varies across the distribution. We do this in Section 5.2.

5.1 The Role of Sex and Education

Figure 6 shows the unadjusted and adjusted wage growth for men and for women. The figure suggests that there is very little difference in the unadjusted wage growth between men and women and that the composition adjustment has very similar effects on both.

Figure 7 shows the unadjusted and adjusted wage growth for individuals with educational attainment of high school or less as well as for individuals with additional education. The main difference in the unadjusted wage growth between the two groups is that the higher educated individuals experienced a faster growth in wages at the onset of the pandemic, and that the lower educated individuals experienced lower average wage increases (and higher wage declines) at the end of the pandemic. The adjusted wage growths are much more similar in the two groups.

The large adjustment in wage growth for the higher educated in 2020 is consistent with Table 1. Conditional on being in the lower education groups (Less than High School and High School degree), the probability of belonging to either of them changed by less than 1.5 percentage points. In contrast, among the higher educated, the probability of being in the group Some College decreased by almost 5 percentage points. In other words, there was a larger change in the composition of workers within the higher educated than within the lower educated. This is consistent with a larger adjustment in the former group (although, of course, the adjustment also depends on characteristic other than education).

5.2 Wage Growth Across the Distribution of Wages

In the discussion above, we used the trimmed mean to calculate wage growth. The trimming was at the 15^{th} and 85^{th} percentiles, which emphasizes a large symmetric "middle part" of the distribution. We can focus on other parts of the distribution by using other percentiles. In Figure 8, we present the unadjusted and composition-adjusted wage growth using the 10^{th}-30^{th}, 40^{th}-60^{th}, and 70^{th}-90^{th} percentiles.

Comparing across the three panels in Figure 8, one sees that the composition adjustment makes a substantial difference between 2020 and 2022, with the adjustments being more important at the middle and upper end of the distribution. Interesting, the figure also indicates that wage growth has been larger at the lower part of the distributions from approximately 2022 until the start of 2024. Figure 9 explores this further by showing the evolution of growth in the trimmed mean of wages for different parts of the distribution. It appears that except for the pandemic, and perhaps the most recent periods, the there has been stronger growth at the lower tail of the distribution. This is consistent with other findings in the literature. See, for example, the discussion in Aaronson, Hu, and Rajan (2020).

6 Sampling Weights

Virtually all survey-based datasets are subject to non-response or attrition, which can lead to non-representative samples—even when the initial sample is randomly drawn from the population. A common solution to this issue is the use of sampling weights, which can be used to adjust the sample to be representative along specific dimensions. For example, weights may be constructed to align the sample's distribution of education or gender with that of the broader population. The Current Population Survey is one such dataset that provides several sampling weights, including one designed to make the weighted distribution of earnings representative of the population's earnings distribution.[6]

In principle, the adjustment methods proposed in this paper can be adapted to incorporate sampling weights. For both the distributional regression and matching approaches, each begins with an observation at time t and estimates what wage that individual would have earned in period $t-1$. Sampling weights from period t can, in theory, be applied in this context. On the other hand, incorporating sampling weights into the reweighting approach of DiNardo, Fortin, and Lemieux (1996) is more complex. This method starts with observations from period $t-1$ and reweights them so their characteristics match those in period t. To apply sampling weights here, one would need to know – or estimate – the sampling weight that an individual observed in $t-1$ would have had in t. Since the CPS does not publish the algorithm used to compute earnings-related sampling weights as a function of observable variables, we do not pursue this approach further.

Figure 10 provides some reassurance that omitting sampling weights does not materially affect our results. The figure shows the unadjusted annual percentage change in the 15–85% trimmed mean of wages, both with and without sampling weights. Although there is a small divergence around 2015, the two series closely track each other, especially during and after the COVID-19 pandemic. This suggests that the adjustments proposed in this paper likely have a much greater impact on wage growth measurement than the application of sampling weights.

[6] It is important to emphasize that sampling weights do not resolve the fundamental selection problem addressed in this paper. In an ideal case, sampling weights correct for representativeness at a single point in time. The issue we focus on here is the change in the distribution of worker characteristics over time.

7 Conclusion

Accurately measuring wage growth requires distinguishing between true increases in pay for a given type of worker and job, and changes that result from shifts in workforce composition or job characteristics. This paper has aimed to develop a wage growth measure that accounts for both demographic changes in the workforce and variation in job types. Our proposed approach accommodates a wide range of observed worker and job characteristics and is designed to be robust to common survey data issues, including heaping, rounding, outliers, and top-coding.

Our findings show that adjusting for observable worker and job characteristics can significantly influence measured wage growth—especially during periods of major labor market disruption, such as the Great Recession and the COVID-19 pandemic. The method we propose enables a more refined analysis of wage dynamics, allowing for meaningful comparisons across demographic groups and segments of the wage distribution.

While our approach focuses on observed characteristics, extending the framework to account for unobserved heterogeneity remains an important direction for future research. Additionally, although our fixed distribution of worker and job characteristics is defined relative to the sample, adapting it to reflect the broader population is feasible using sampling weights and may enhance its relevance for policy applications.

Overall, the methodology introduced in this paper provides a flexible and robust tool for analyzing wage growth. By addressing key limitations in existing measures, it offers a valuable foundation for future work on labor market dynamics and supports more informed economic policymaking.

References

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

Aaronson, D., and D. G. Sullivan (2001): "Growth in worker quality," Economic Perspectives, 25(4).

Blinder, A. S. (1973): "Wage Discrimination: Reduced Form and Structural Estimates," The Journal of Human Resources, 8(4), 436–455.

Chernozhukov, V., I. Fernández-Val, and B. Melly (2013): "Inference on Counterfactual Distributions," Econometrica, 81(6), 2205–2268.

Cole, R., L. Hu, and S. Schulhofer-Wohl (2017): "When It Comes to Wage Growth, the Measure Matters," Chicago Fed Letter, (387).

DiNardo, J., N. M. Fortin, and T. Lemieux (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," Econometrica, 64(5), 1001–1044.

Escanciano, J. C., D. Jacho-Chávez, and A. Lewbel (2016): "Identification and estimation of semiparametric two-step models," Quantitative Economics, 7(2), 561–589.

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, D. Backman, A. Chen, G. Cooper, S. Richards, M. Schouweiler, and M. Westberry (2024): "IPUMS CPS: Version 12.0 [dataset]," .

Flood, S., M. King, R. Rodgers, S. Ruggles, J. R. Warren, and M. Westberry (2021): "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]," Minneapolis, MN: IPUMS, 2021.

Foresi, S., and F. Peracchi (1995): "The Conditional Distribution of Excess Returns: An Empirical Analysis," Journal of the American Statistical Association, 90(430), 451–466.

FORTIN, N. M., T. LEMIEUX, AND S. FIRPO (2010): "Decomposition Methods in Economics," in Handbook of Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 4A, pp. 1–102. Elsevier.

HECKMAN, J. J., H. ICHIMURA, AND P. E. TODD (1997): "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," The Review of Economic Studies, 64(4), 605–654.

——— (1998): "Matching as an Econometric Evaluation Estimator," The Review of Economic Studies, 65(2), 261–294.

HONORÉ, B. E., AND L. HU (2020): "Selection Without Exclusion," Econometrica, 88, 1007–1029.

——— (2023): "The COVID-19 pandemic and Asian American employment," Empirical Economics, 64(5), 2053–2083.

IMBENS, G. W., AND J. M. WOOLDRIDGE (2009a): "Recent Developments in the Econometrics of Program Evaluation," Journal of Economic Literature, 47(1), 5–86.

——— (2009b): "Recent Developments in the Econometrics of Program Evaluation," Journal of Economic Literature, 47(1), 5–86.

KITAGAWA, E. M. (1955): "Components of a Difference Between Two Rates*," Journal of the American Statistical Association, 50(272), 1168–1194.

OAXACA, R. (1973): "Male-Female Wage Differentials in Urban Labor Markets," International Economic Review, 14(3), 693–709.

OAXACA, R. L., AND E. SIERMINSKA (2023): "Oaxaca-Blinder Meets Kitagawa: What Is the Link?," Discussion paper, IZA - Institute of Labor Economics.

RINZ, K., AND J. VOORHEIS (2018): "The Distributional Impact of the Minimum Wage in the Short and Long Run," Discussion Paper CES-WP-18-49, U.S. Census Bureau Center for Economic Studies.

Rosenbaum, P. R., and D. B. Rubin (1983): "The central role of the propensity score in observational studies for causal effects," Biometrika, 70(1), 41–55.

Tobin, J. (1958): "Estimation of Relationships for Limited Dependent Variables," Econometrica, 26(1), 24–36.

Data Appendix

The data used in this paper are based on the Basic Monthly Current Population Survey (CPS). We use the version provided by IPUMS at https://cps.ipums.org/cps/ (see Flood, King, Rodgers, Ruggles, Warren, and Westberry (2021)) and the sample covers the period from January 2006 to December 2024.[7] We strict the data to individuals aged 16 to 70 (inclusive) and we exclude military personnel.

Most of the variables are used "as is", but a few require explanation.

- The variable citizen is a dummy that is one if the individial was born in the U.S. (including outlying), born abroad of American parents, or is a naturalized citizen.
- The variable occupation is based on occ2010 aggregated into the intervals $[10, 1000)$, $[1000, 3700)$, $[3700, 4700)$, $[4700, 5000)$, $[5000, 6000)$, $[6000, 7000)$, $[7000, 7700)$, $[7700, 9000)$, $[9000, 9800)$, $[9800, 9999]$. (See https://cps.ipums.org/for description)
- The variable industry is based on ind1990 aggregated into the intervals $(-\infty, 10)$, $[0, 10)$, $[10, 40)$, $[40, 60)$, $[60, 100)$, $[100, 400)$, $[400, 500)$, $[500, 700)$, $[700, 721)$, $[721, 761)$, $[761, 800)$, $[800, 812)$, $[812, 900)$, $[900, 940)$, $[940, 998]$. (See https://cps.ipums.org/for description)
- The education variables are no high school degree, high school degree or equivalent, four year college degree, and more.
- The wage rate is constructed as follows. We start with the variable hourwage2. If this variable is missing or allocated, then it is the ratio of weekly earnings (earnweek2) to the usual hours worked per week (uhours) (unless either of these variables is missing or allocated)
- In the construction of race/ethnicity Hispanic refers to Hispanic White.

[7]The data were downloaded March 26, 2025 and reflect IPUMS's harmonization as of that date. Accessing the data through IPUMS requires registration and general redistribution is not allowed.

Figure 1: Comparison of Common Measures of Annual Wage Growth

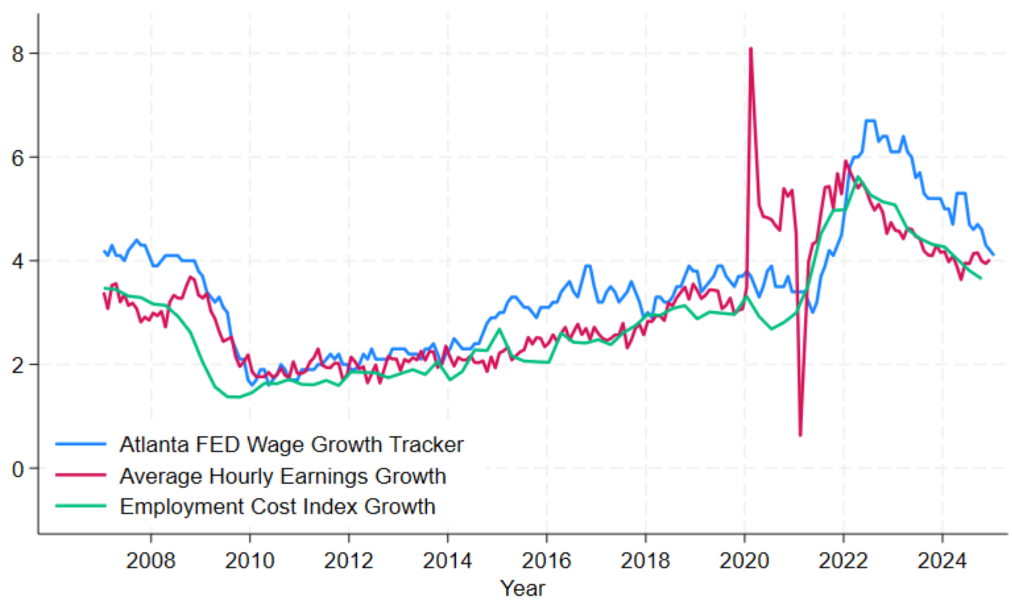

Note: The figure plots three common measures of annual wage growth. The Atlanta Fed's Wage Growth Tracker was downloaded from the Federal Reserve Bank of Altanta website while the percentage change in the Bureau of Labor Statistics' Average Hourly Earnings (Average Hourly Earnings of All Employees, Total Private), and the percentage change in the Bureau of Labor Statistics' Employment Cost Index (Employment Cost Index: Wages and Salaries: Private Industry Workers) were downloaded from FRED (Federal Reserve Bank of St. Louis).

Figure 2: Raw and Composition-Adjusted Percentage Growth in Wages

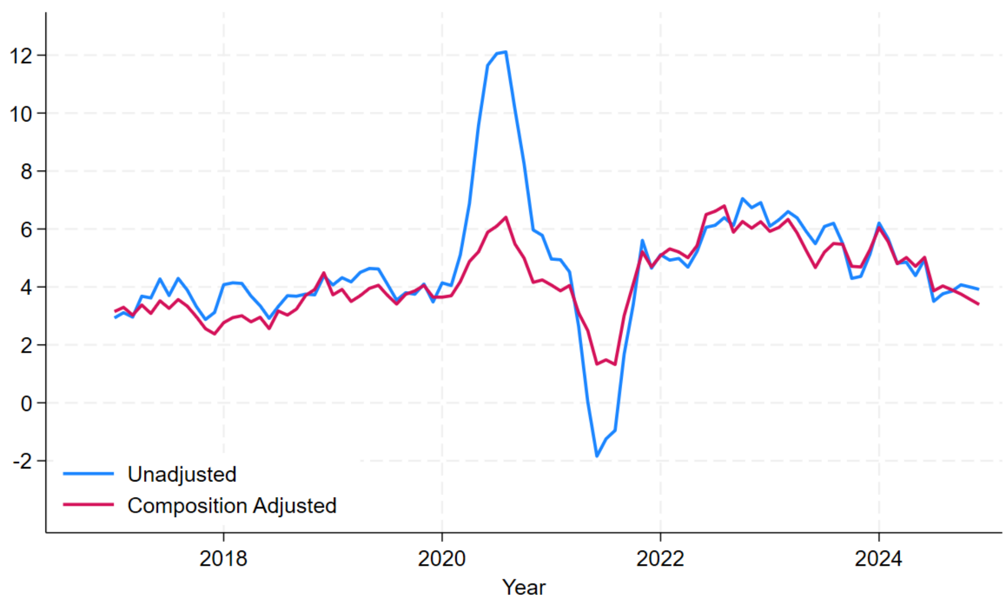

Note: The figure plots the annual percentage growth in the 15%-85% trimmed mean of the hourly wage. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2007 to December 2024.

Figure 3: Percentage Change in 15-85 Trimmed Mean Compared to Change in Overall Mean

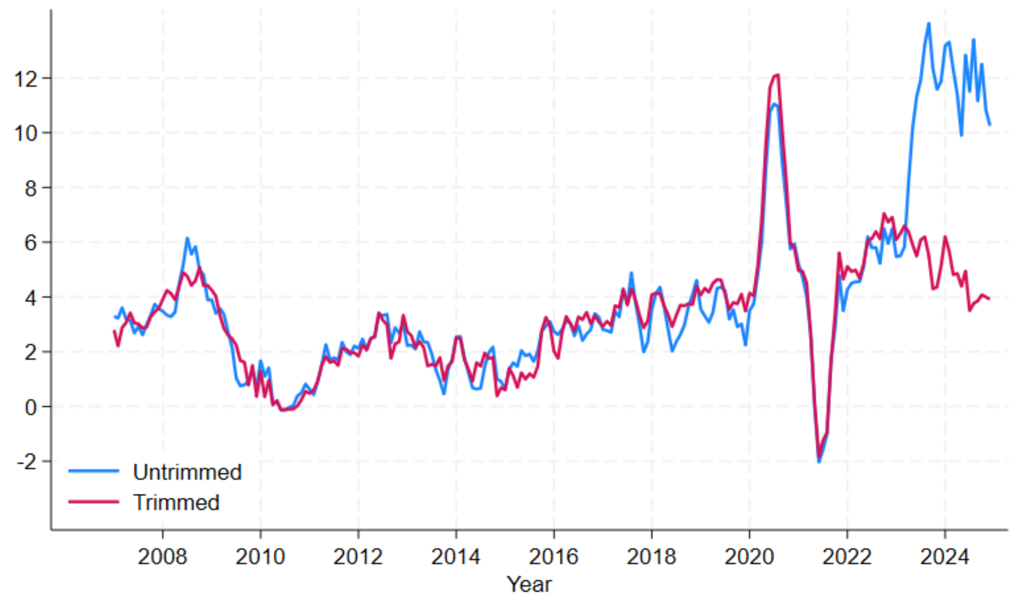

Note: The figure plots the annual percentage growth in the mean and overall 15%-85% trimmed mean of the hourly wage. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2017 to December 2024.

Figure 4: Composition-Adjusted Wage Growth (Matching, Reweighting and Dist. Reg.)

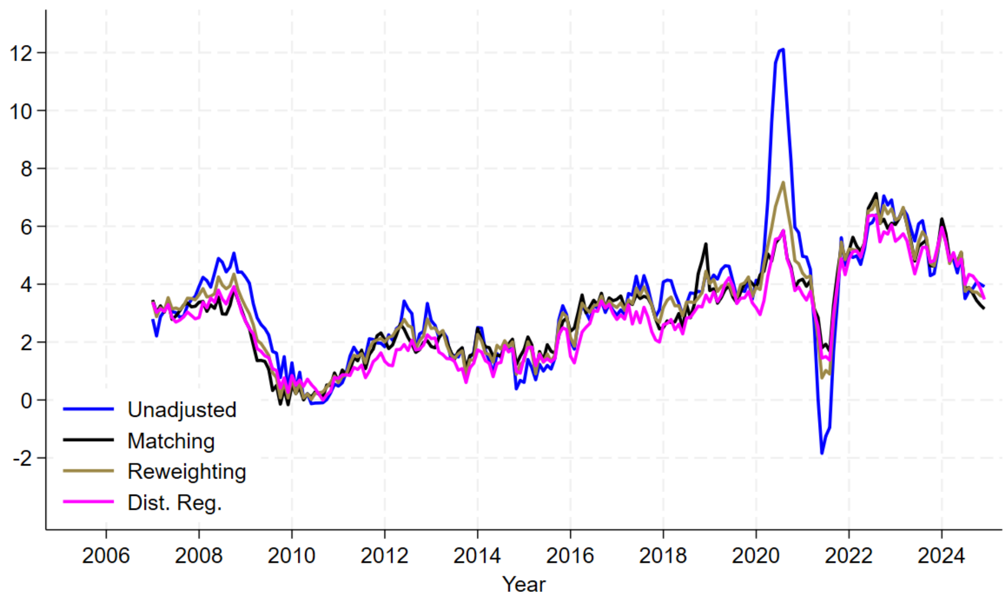

Note: The figure plots the annual percentage growth in raw and adjusted 15%-85% trimmed mean of the hourly wage using each of the three adjustment methods. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2007 to December 2024.

Figure 5: Raw and Composition-Adjusted Percentage Growth in Wages

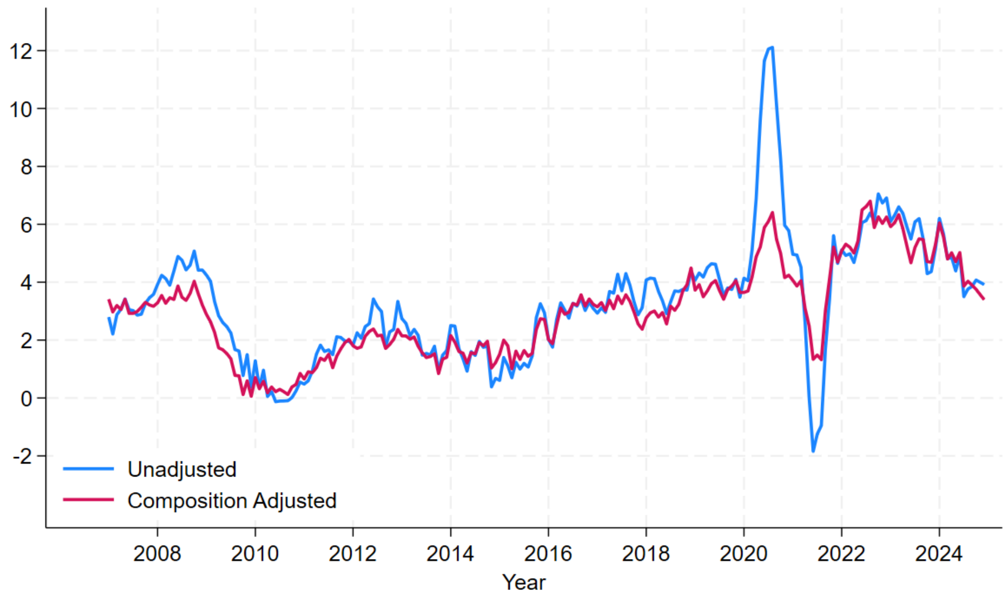

Note: The figure plots the annual percentage growth in raw and adjusted 15%-85% trimmed mean of the hourly wage. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2007 to December 2024.

Figure 6: Unadjusted and Adjusted Wage Growth For Women and For Men

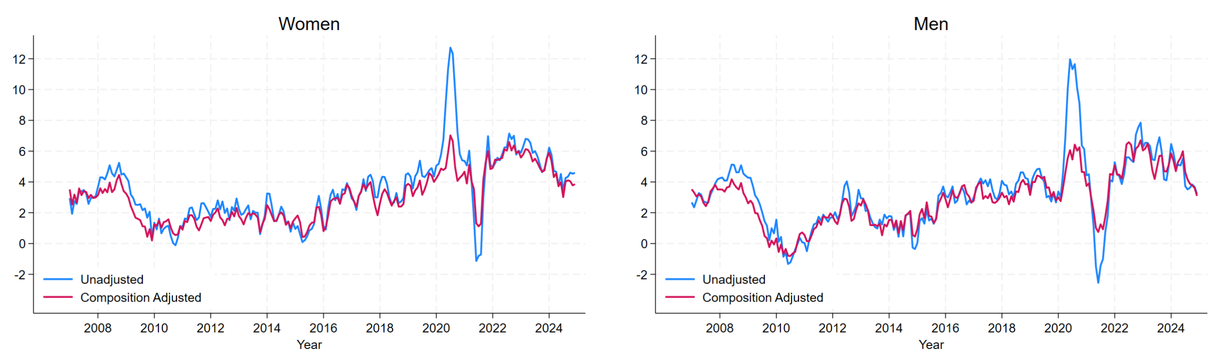

Note: The figures plot the annual percentage growth in raw and adjusted 15%-85% trimmed mean of the hourly wage for women (top panel) and for men (bottom panel). The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2007 to December 2024.

Figure 7: Unadjusted and Adjusted Wage Growth By Education

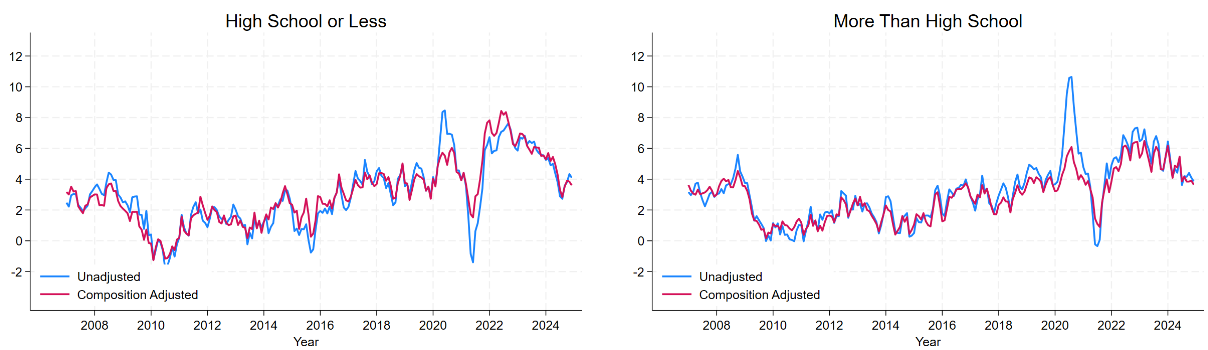

Note: The figures plot the annual percentage growth in raw and adjusted 15%-85% trimmed mean of the hourly wage. The top panel uses data for individual with a high school degree or less while the bottom panel uses data for individuals with more than a high school degree. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2017 to December 2024.

Figure 8: Wage Growth By Percentiles

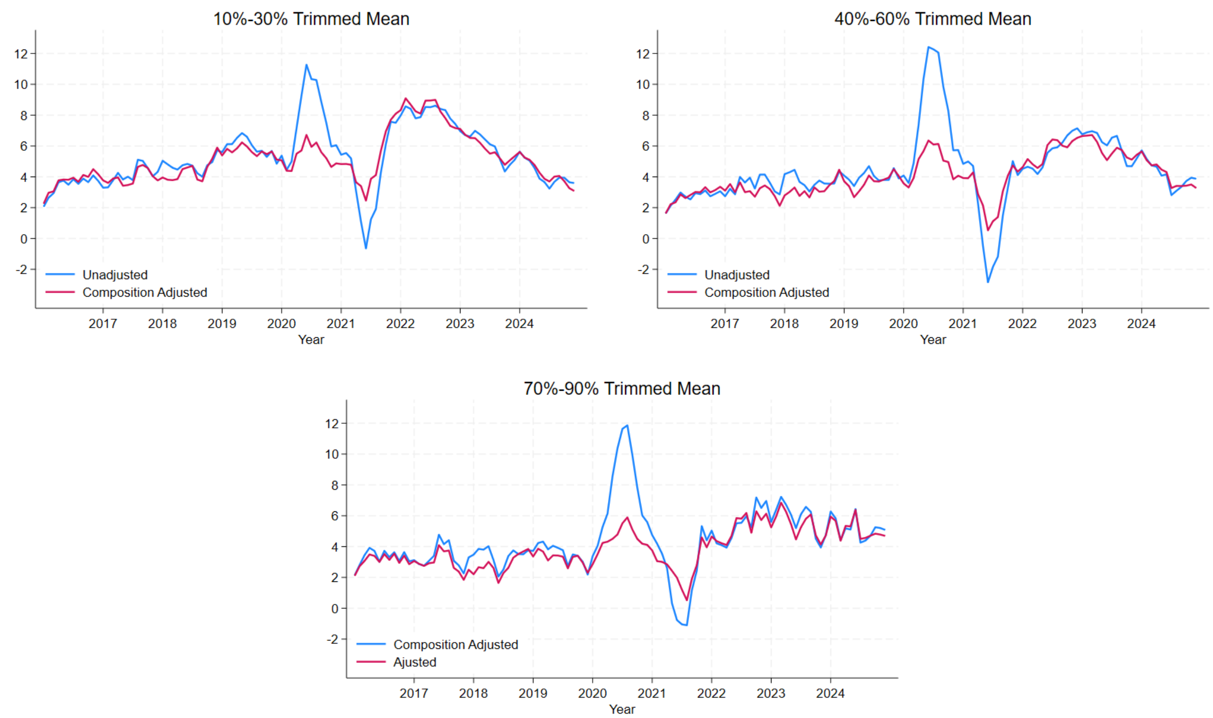

Note: The figures plot the annual percentage growth in raw and adjusted 10%-30%, 40%-60% and 70%-90%trimmed means of the hourly wage. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2017 to December 2024.

Figure 9: Wage Growth Across the Wage Distribution

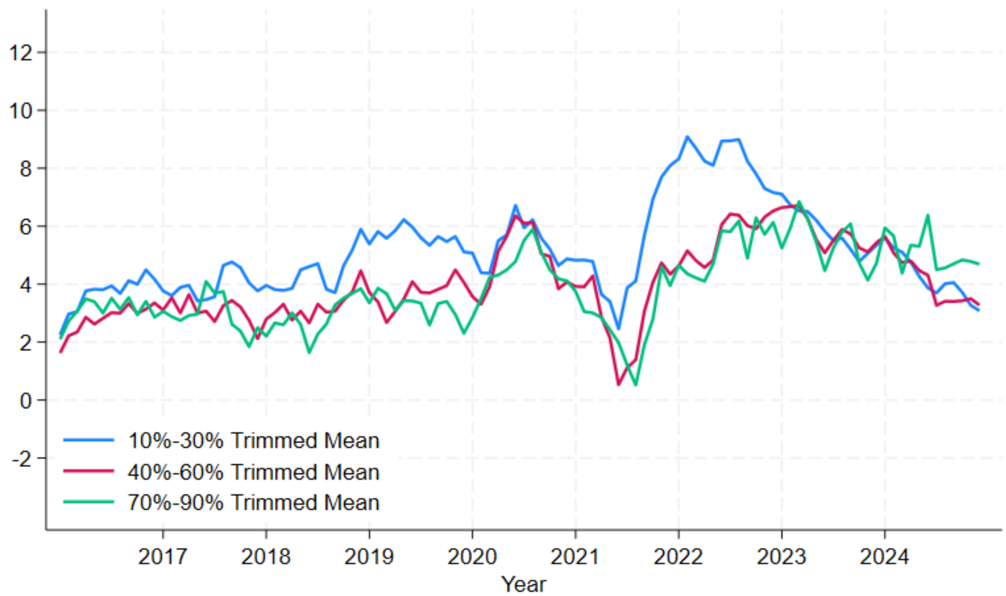

Note: The figures plot the annual percentage growth of adjusted 10%-30%, 40%-60% and 70%-90 trimmed means of the hourly wage. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2017 to December 2024.

Figure 10: Raw Change in 15-85 Trimmed Mean With and Without Sampling Weights

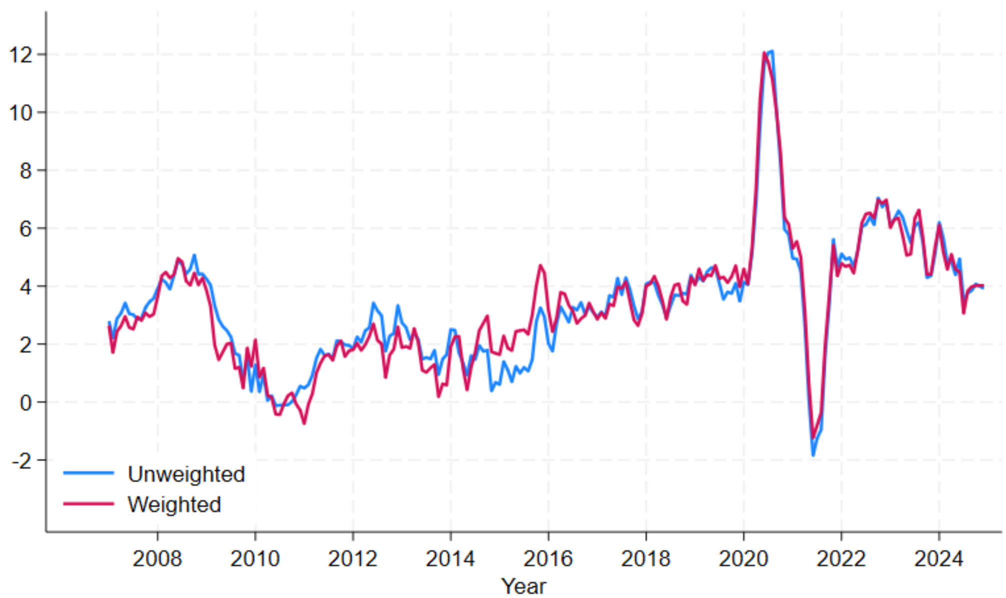

Note: The figure plots the annual percentage growth in the 15%-85% trimmed mean of the hourly wage with and without sampling weights. The data is from IPUMS CPS, and the graph covers individuals aged 16-70 over the period from January 2017 to December 2024.