STAFF
REPORTS

NO. 1067
JULY 2023

REVISED
NOVEMBER 2024

A Measure of Trend Wage Inflation

Martín Almuzara | Richard Audoly | Davide Melcangi

FEDERAL RESERVE BANK *of* NEW YORK

A Measure of Trend Wage Inflation
Martín Almuzara, Richard Audoly, and Davide Melcangi
Federal Reserve Bank of New York Staff Reports, no. 1067
July 2023; revised November 2024
JEL classification: C33, E24

Abstract

We extend time-series models that have so far been used to study price inflation (Stock and Watson [2016a]) and apply them to a micro-level dataset containing worker-level information on hourly wages. We construct a measure of aggregate nominal wage growth that (i) filters out noise and very transitory movements, (ii) quantifies the importance of idiosyncratic factors for aggregate wage dynamics, and (iii) strongly co-moves with labor market tightness, unlike existing indicators of wage inflation. We show that our measure is a reliable real-time indicator of wage pressures and a good predictor of future wage growth.

Key words: wage inflation, persistence, factor models

Almuzara, Audoly, Melcangi: Federal Reserve Bank of New York (emails: martin.almuzara@ny.frb.org, richard.audoly@ny.frb.org, davide.melcangi@ny.frb.org). This paper was previously circulated under the title “A Measure of Core Wage Inflation.” The authors thank Mary Amiti, Richard Crump, Marco Del Negro, Keshav Dogra, Danial Lashkari, Silvia Sarpietro, Argia Sbordone, Giorgio Topa, and John C. Williams for helpful comments and discussions. They also thank Augustin Belin, Julie Seager, Aidan Toner-Rodgers, and Roshie Xing for excellent research assistance.

1 Introduction

The recent surge in inflation has renewed interest in the aggregate evolution of nominal wages. Because wage inflation is perceived as tightly linked to price inflation, it is one of the key indicators monitored for the conduct of monetary policy. It also provides a signal about the state of the labor market, and it is an important input for households' and firms' decisions. While there is extensive work on constructing measures of price inflation that are purged from noise and short-run fluctuations, surprisingly little research has been devoted to the underlying dynamics of aggregate nominal wage growth.[1]

This paper describes a framework to isolate the latent persistent ("trend") component of wage inflation. We estimate a dynamic factor model with time-varying parameters on worker-level wage data. Aggregate wage inflation is the sum of a persistent component common to all workers, a persistent component specific to some subgroup of workers, such as industries or occupations, and transitory shocks. By filtering out noise and very transitory movements, our measure of Trend Wage Inflation strongly co-moves with indicators of labor market tightness, unlike existing indicators of wage inflation, and our measure therefore represents a reliable real-time indicator of wage pressures. Moreover, our framework sheds light on the drivers of aggregate wage inflation. We find that business cycle fluctuations in wage inflation are primarily driven by the persistent component that is common across workers. This common latent factor is especially predominant during inflationary episodes, such as the one of 2021.

Our analysis employs monthly worker-level information on hourly wages from the Current Population Survey (CPS). Our statistical model lends itself naturally to this data set since it allows us to break down aggregate changes in nominal wages across many alternative cross-sections of the data. In line with the price inflation literature, our main

[1]See Stock and Watson [2016a] for a summary of the price inflation literature. On the evolution of aggregate wages, Daly et al. [2011] examine the time-varying cyclicality of real wages. Huh and Trehan [1995] estimate a vector error correction model with prices, wages and productivity. A broader literature studies the link between nominal wage growth and price inflation, which is beyond the scope of our paper. For recent work on this, see Kiley [2023].

focus is on the industry cross-section, but we also consider a range of alternatives based on worker and job characteristics. A key difference of this data set relative to price inflation data is that, although observed each month, nominal wage growth in our worker-level data is measured as a 12-month rate of change instead of as a month-on-month growth rate.

The *first contribution* of our paper is to design a model that explicitly accounts for the monthly-to-yearly temporal aggregation. We build on the framework that Stock and Watson [2016a] previously applied to price data to retrieve the persistent component of unobserved monthly wage growth.[2] This component is our key object of interest, and we refer to it as Trend Wage Inflation (TWIn hereafter). To validate our approach, we conduct a real-time forecast comparison against a random-walk benchmark. Random-walk forecasts are known to be hard to improve upon in the price inflation literature (see, e.g., Atkeson and Ohanian [2001], Stock and Watson [2008])). Yet we find that our approach outperforms the random-walk alternative, especially at short horizons and at the industry level.

A *second contribution* of our paper is to empirically document the dynamics of wage inflation and its drivers using our model. While wage growth is characterized by substantial monthly variation, the largest movements in TWIn coincide with the episodes of greatest macroeconomic significance in our sample. TWIn falls substantially during the 2001 and 2008 recessions, and rises sharply at the beginning of 2021, at the onset of the post-pandemic inflation surge episode.

We quantify the relative importance of economy-wide versus industry-specific shocks for the evolution of aggregate wage inflation. We find that most of the variation in the persistent component of nominal wage growth is common across industries, especially during the post-pandemic surge in inflation. By contrast, the estimated sector–specific

[2] This methodological contribution is novel and potentially relevant for the empirical analysis of other data sets subject to temporal aggregation. Importantly, we developed an efficient implementation of our framework that can be easily replicated and used in other applications.

component appears to capture long-run trends, and it has been declining since the 1990s, albeit tentatively picking up since 2020.

The dynamics of aggregate nominal wage growth can be seen as providing a signal of the state of the labor market, and the extent of labor market imbalances. To this end, it is important to evaluate how synchronized nominal wage growth is with labor market tightness and measures of price inflation. We find that changes in TWIn (and in its common component) are strongly contemporaneously correlated with changes in the vacancy-to-labor force ratio, the unemployment rate, and measures of price inflation. These correlations are stronger than for other popular measures of aggregate wage growth. As a result, our estimates tend to lead these alternative measures and appear better aligned with labor market tightness.[3]

As such, we see it as a *key contribution* of our paper to construct a timely indicator to assess the state of the labor market in real time. This is especially valuable when macroeconomic conditions evolve rapidly. The recent COVID-19 pandemic has shown the importance of having reliable, timely, indicators of economic conditions [see, for instance, Lewis et al., 2022]. Successful indicators do not only need to be available at high frequency, but also need to look through noisy fluctuations, while extracting the signal from substantial economic movements. Our approach does so by using monthly data, by filtering these data appropriately to extract the persistent part of wage growth, and by formally accounting for time aggregation.

Our *last contribution* is to document the role of worker heterogeneity for aggregate nominal wage growth. We do this by re-estimating our model using alternative cross-sections of the data based on race, gender, education, age, and other demographics. Our findings about the dynamics of wage inflation and the synchronicity with labor market tightness are robust. Most importantly, none of the alternative cross-sections affects the

[3]We focus on the contemporaneous correlation between measures of changes in labor market tightness and measures of wage growth at the monthly frequency. We think that TWIn can also be relevant for variables at different frequencies, such as GDP growth or ECI. This can be achieved by means of mixed data sampling techniques (see, Ghysels [2012] and Ghysels et al. [2020] for comprehensive surveys).

conclusion that aggregate nominal wage growth – and, in turn, TWIn – is mostly driven by its common component during the three episodes of greatest macroeconomic significance in our sample: the 2001 and 2008 recessions, and the post-pandemic inflation surge. To be more specific, the common factor explains a large share of the fall in TWIn during the two recessions (for example, between 50 percent and 85 percent of the 2008-2009 drop across the cross-sections we consider). The pattern is even stronger during the 2021 inflation surge, when the common component accounts for at least 80 percent of the increase in TWIn. Additional estimates using an alternative data set going back to the 1960s suggest that this is not specific to the post-pandemic period: a common, persistent, factor lies behind aggregate nominal wage growth during the inflationary spikes of the 1970s. We see these empirical patterns as a useful benchmark for models of wage determination with aggregate shocks and cross-sectional heterogeneity.

Our work primarily relates to the literature using dynamic factor models to study the evolution of prices. While extensive work has been done on inflation (see Stock and Watson [2016b]), much less research has focused on nominal wage growth. Two exceptions are worth highlighting. Consistent with our empirical findings, Watson and Engle [1981] find that a highly persistent single common factor explains most of the variation across industries using wage data for the metropolitan area of Los Angeles. In recent work, Ahn et al. [2024] use a dynamic factor model on quarterly data from the Current Employment Survey (CES) to extract the common component of inflation in average hourly earnings across a large number of industries. In addition to the data, an important difference is that Trend Wage Inflation measures both the common and sector-specific persistent components of wage inflation while Ahn et al. [2024] retrieve the common persistent component of wage inflation. Some of our empirical findings are consistent with Ahn et al. [2024]: most notably, both measures are more strongly correlated with movements in labor market tightness than other measures of nominal wage growth. This suggests that dynamic factor models can be useful to look through short-run fluctuations and isolate fluctuations in nominal wage growth that are more tightly associated with labor market and price setting patterns.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 describes the model and estimation method. Section 4 develops our empirical analysis, and Section 5 concludes.

2 Wage inflation data

Figure 1 shows three commonly used measures of wage inflation derived from US data: the Atlanta Fed Wage Growth Tracker (AWGT), Average Hourly Earnings (AHE), and the Employment Cost Index (ECI). We make three observations. First, composition matters. The AWGT and the ECI adjust for changes in the composition of the workforce by considering, respectively, the median wage growth within-worker and the average compensation growth within-job. By contrast, wage inflation based on AHE is growth in average hourly earnings. As a result, AHE exhibits large swings following the onset of the Covid pandemic, when restrictions affected many low-paid jobs. These movements are not present in the AWGT and the ECI. Similarly, growth in AHE slightly increases during the Great Recession (2007m12-2009m6), while the AWGT and growth in the ECI start decreasing during the recession. Second, the frequency at which these series are released differ across data sets. AHE and the AWGT are released monthly, while the ECI is quarterly. Third, the AWGT and, to a lesser extent, AHE tend to be volatile in the short-run.

A key objective of this paper is to propose a measure of wage inflation that provides a reliable signal for the evolution of the marginal cost of labor in real time. We then need to (i) account for changes in the composition of the workforce, and (ii) use data released on a regular basis. In our baseline analysis, we therefore use monthly data on nominal wage

FIGURE 1. Measures of Wage Inflation in the US

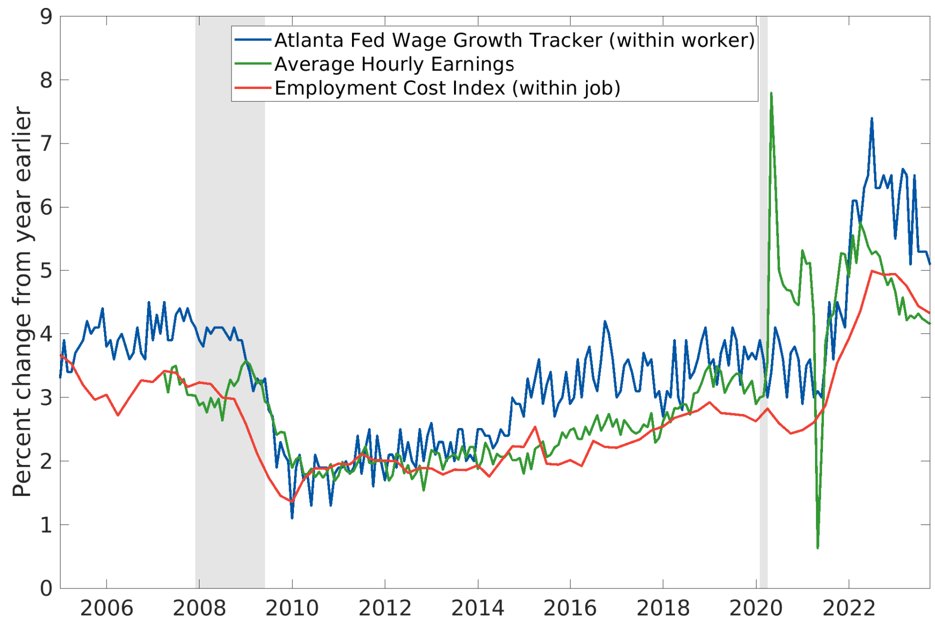

NOTES. Measures of the 12-month rate of wage inflation over the period 2005m1-2023m9.

growth at the worker level from the Current Population Survey (CPS).[4] Specifically, we define wage inflation as the median growth in the hourly wage of individuals observed twelve months apart.[5] Letting t denote a month and i denote some partition of the data, such as industries or occupation groups, the 12-month rate of wage inflation $w_{i,t}$ in i is defined as

$$w_{i,t} = \ln\left(\frac{W_{i,j50,t}}{W_{i,j50,t-12}}\right), \tag{1}$$

where $j50$ denotes the individual with median 12-month wage growth in group i and W denotes hourly wages.

[4]We use data between 1997m1 and 2023m9. Our measure can be updated each month after the CPS monthly basic files are released.

[5]We show in Appendix E that our results are similar, albeit with less precise estimates, if we use average growth. The hourly wage is for the person's "usual" weekly earnings and "usual" hours, which include the variable part of compensation (overtime pay, tips, commissions). This definition is similar to the Atlanta Fed WGT, though the headline figure reported by the Atlanta Fed (Figure 1) is a three-month moving average [Atlanta Fed, 2023].

We consider individual wage changes from observations twelve months apart due to the interview structure of the CPS.[6] This directly adjusts for seasonal factors. But it also implies that this measure of wage inflation tends to lag actual month-on-month changes. To see this, consider the case of a worker with no month-on-month wage growth for eleven months and 0.1 percent in the twelfth month. Their twelve-month apart wage change is 0.1 percent, but their annualized month-on-month change in wages, the object of interest to monitor wage inflation, is $12 \times 0.1\% = 1.2\%$. We fully account for time aggregation in the measurement framework introduced in Section 3.

An attractive feature of building our measure from worker-level data is that aggregate wage inflation can be decomposed by both job and demographic characteristics. For most of Section 4, we break down aggregate changes in nominal wages by industry. We consider seven broad industry groups, which ensures that the sample size within each cell is not too small: Construction and Mining, Education and Health, Finance and Business Services, Leisure and Hospitality, Manufacturing, Public Administration, and Trade and Transportation. This classification refers to workers' current industry, irrespective of their industry in their previous interview. In Section 4.3, we consider many alternative worker characteristics, such as occupation groups, education groups, regions, and wage quartiles. The definitions of these variables are given in Appendix A.

By defining wage inflation as wage changes for the same workers, our CPS measure can be seen as the worker equivalent to standard measures of price inflation, which summarize price changes for the same basket of goods and services. The ECI also follows a similar approach by considering wage changes within jobs.[7] While within-worker and within-job wage changes in part adjust for composition effects (see Figure 1 and the literature on wage cyclicality using worker-level data starting with Bils [1985]), there is still potentially selection in the type of workers who switch jobs and the type of jobs

[6]Individuals are interviewed for four consecutive months, spend eight month out of the sample, and are subsquently interviewed again for four additional months.

[7]A job in the ECI is a small sector-industry-occupation cell. See the BLS website for definitions: https://www.bls.gov/opub/hom/ncs/home.htm.

available to these workers at any point in time. An active literature studies the cyclicality of wages attempting to control for worker selection in job switching [Gertler et al., 2020, Grigsby et al., 2021, Hazell and Taska, 2020]. In line with the well-established definition underlying the AWGT, we do not impose further restrictions on the sample of workers. While we use CPS data as our baseline, we also stress that the econometric framework we outline in the next section can be estimated for any definition of wage inflation.

3 Model and estimation

We describe our modeling approach and our proposed measure of Trend Wage Inflation in this section. A challenge posed by the data we analyze is that nominal wage growth, although observed every month, is measured as a 12-month growth rate. We are interested in the persistence of monthly wage growth, and we therefore need to explicitly account for the monthly-to-yearly temporal aggregation.

Let w_{it} be the 12-month rate of wage inflation for sector i and month t as defined in Equation (1).[8] We decompose w_{it} into the contribution of a persistent (or trend) component and a transitory (or noise) component:

$$w_{it} = \frac{1}{12} \sum_{\ell=1}^{12} \tilde{\tau}_{i,t+1-\ell} + \tilde{\varepsilon}_{it}. \tag{2}$$

In Equation (2), $\tilde{\tau}_{it}$ is the persistent component for the unobserved monthly (annualized) rate of wage growth, while $\tilde{\varepsilon}_{it}$ is the transitory error in the observed 12-month rate. To capture the potential cross-sectional correlation across sectors, we further decompose trend and noise into the sum of common (indexed by c) and sector-specific components:

$$\tilde{\tau}_{it} = \alpha_{\tau,it} \tau_{ct} + \tau_{it},$$

[8] Here, "sectors" generically denote any worker-level partition we consider, such as industry (our baseline), age, occupations or regions.

$$\tilde{\varepsilon}_{it} = \alpha_{\varepsilon,it}\varepsilon_{ct} + \varepsilon_{it},$$

and we model trends as random walks and transitory errors as low-order moving averages,

$$\begin{aligned}
\tau_{ct} &= \tau_{c,t-1} + \sigma_{\Delta\tau,ct}\eta_{\Delta\tau,ct},\\
\tau_{it} &= \tau_{i,t-1} + \sigma_{\Delta\tau,it}\eta_{\Delta\tau,it},\\
\varepsilon_{ct} &= (1 + \theta_{c1}L + \cdots + \theta_{cp}L^{p})\sigma_{\varepsilon,ct}\eta_{\varepsilon,ct},\\
\varepsilon_{it} &= (1 + \theta_{i1}L + \cdots + \theta_{iq}L^{q})\sigma_{\varepsilon,it}\eta_{\varepsilon,it}.
\end{aligned}$$

Here, L is the lag operator and $\eta_{\Delta\tau,ct}, \eta_{\Delta\tau,it}, \eta_{\varepsilon,ct}, \eta_{\varepsilon,it}$ are serially and mutually independent $N(0,1)$, and independent of the loadings $\{\alpha_{\tau,it}, \alpha_{\varepsilon,it}\}_{i=1}^{n}$ and volatilities $\sigma_{\Delta\tau,ct}$, $\{\sigma_{\Delta\tau,it}\}_{i=1}^{n}$, $\sigma_{\varepsilon,ct}$, $\{\sigma_{\varepsilon,it}\}_{i=1}^{n}$.

Time aggregation The monthly-to-yearly transformation is an important feature of our framework and a key difference with models of price inflation such as Stock and Watson [2016a] in which monthly rates of change are observed directly. To interpret what we do, rewrite the definition of wage inflation (1) as

$$w_{it} = \ln\left(\frac{W_{i,j50,t}}{W_{i,j50,t-12}}\right) = \ln\left(\frac{W_{i,j50,t}}{W_{i,j50,t-1}} \times \frac{W_{i,j50,t-1}}{W_{i,j50,t-2}} \times \cdots \times \frac{W_{i,j50,t-11}}{W_{i,j50,t-12}}\right),$$

where $W_{i,j50,t-1}, \ldots, W_{i,j50,t-11}$ are unobserved because of the sampling design of the CPS. Consider the case $\tilde{\varepsilon}_{it} = 0$ with no noise component in (2). Then,

$$\tilde{\tau}_{it} = 12\ln\left(\frac{W_{i,j50,t}}{W_{i,j50,t-1}}\right) \approx \left(\frac{W_{i,j50,t} - W_{i,j50,t-1}}{W_{i,j50,t-1}}\right)^{12},$$

that is, $\tilde{\tau}_{it}$ would measure the unobserved monthly annualized growth in median wage for worker $j50$.

In practice, w_{it} is an estimate based on survey data and expected to contain measurement error. Due to the sampling design of the CPS, the workers who contribute to $W_{i,j50,t}/W_{i,j50,t-12}$ and to $W_{i,j50,s}/W_{i,j50,s-12}$ differ whenever $t \neq s$. If workers are sampled at random, the measurement error in w_{it} and w_{is} will then be independent for $t \neq s$. The measurement error in w_{it} will be independent of that in w_{jt} for $j \neq i$ by a similar argument. This suggests having no common transitory error and white noise dynamics for the sector-specific transitory components. For added flexibility, we do include ε_{ct} and we allow ε_{it} to be serially dependent through $q > 0$ and the presence of time-varying volatility. Consistent with interpreting $\tilde{\varepsilon}_{it}$ as measurement error, we find ε_{ct} to be small and $\theta_{i\ell}$ to be close to zero. Importantly, $\sigma_{\varepsilon,it}$ appears to evolve in proportion to the inverse of the square root of the number of workers surveyed each month, which is exactly what one would expect if $\tilde{\varepsilon}_{it}$ represented the error in estimating median wage growth in the population using its sample counterpart.[9]

What about genuine transitory shocks to monthly wage growth? To accommodate them we would either enrich $\tilde{\tau}_{it}$ to be itself the sum of persistent and transitory parts or model $\tilde{\varepsilon}_{it}$ as an MA(12). In Appendix E we show that the estimates from this extended model are similar to the more parsimonious model with random walk τ_{ct}, $\{\tau_{it}\}_{i=1}^{n}$ and low p, q. Given the robustness of our results, the empirical analysis reported in Section 4 is based on the model introduced above.

Objects of interest In addition to the sector-level trend $\tilde{\tau}_{it}$ we are interested in the aggregate wage growth trend that we define for sectoral shares $\{s_{it}\}_{i=1}^{n}$ as

$$\tilde{\tau}_t = \sum_{i=1}^{n} s_{it}\tilde{\tau}_{it} = \left(\sum_{i=1}^{n} s_{it}\alpha_{\tau,it}\right)\tau_{ct} + \left(\sum_{i=1}^{n} s_{it}\tau_{it}\right), \tag{3}$$

[9] To be more precise, if n_t is the number of workers surveyed during month t and $\tilde{\sigma}_{\varepsilon,t}$ is the (posterior median estimate of the time-varying) standard deviation of the weighted-average transitory error $\tilde{\varepsilon}_t = \sum_{i=1}^{n} s_{it}\tilde{\varepsilon}_{it}$, we find that most pairs $(n_t, \tilde{\sigma}_{\varepsilon,t})$ lie remarkably close to a line with no intercept, $\tilde{\sigma}_{\varepsilon,t} = 14.1/\sqrt{n_t}$, and that $\text{corr}\left(1/\sqrt{n_t}, \tilde{\sigma}_{\varepsilon,t}\right) = 0.9$. See Appendix E for further discussion.

where the sectoral share s_{it} is defined as the employment share of sector i in period t.[10] We stress that the sectoral shares $\{s_{it}\}_{i=1}^{n}$ are only used to aggregate up the sector-level estimates into a measure of wage inflation for the whole economy. They do not affect the sector-level estimates. We thereafter refer to $\tilde{\tau}_t$ as Trend Wage Inflation (TWIn). Similar to sector-level trends, TWIn is driven both by common persistence across sectors and by a weighted average of sector-specific trend movements. We report this decomposition in Section 4. A technical point to note is that the contribution of common and idiosyncratic components to the *level* of the trend is only pinned down by an arbitrary normalization.[11] The contributions of common and idiosyncratic components to the *changes* in the trend, on the other hand, are not subject to the same caveat. Hence, we focus on changes as opposed to levels when discussing this decomposition below.

Time-varying parameters In specifying the dynamics of the loadings and (log) volatilities, we follow Del Negro and Otrok [2008] and Stock and Watson [2016a] and model them as random walks with small variances, i.e.,

$$\Delta\alpha_{m,it} = \gamma_{\alpha,m,i}\nu_{\alpha,m,it}, \qquad m = \tau, \varepsilon, \quad i = 1, \ldots, n,$$
$$\Delta \ln \sigma^2_{m,jt} = \gamma_{\sigma,m,j}\nu_{\sigma,m,jt}, \qquad m = \Delta\tau, \varepsilon, \quad j = c, 1, \ldots, n,$$

with innovations $\{\nu_{\alpha,\tau,it}, \nu_{\alpha,\varepsilon,it}\}_{i=1}^{n}, \nu_{\sigma,\Delta\tau,ct}, \{\nu_{\sigma,\Delta\tau,it}\}_{i=1}^{n}, \nu_{\sigma,\varepsilon,ct}, \{\nu_{\sigma,\varepsilon,it}\}_{i=1}^{n}$ assumed serially and cross-sectionally independent $N(0,1)$. This is a standard approach to allow for parameters that drift slowly over time. It can accommodate trends in the correlations across sectors and noisiness of the different series.

[10] We use the survey weights included with the CPS to derive the aggregate employment count in each sector i.

[11] This is because the locations of $\{\alpha_{\tau,it}\tau_{ct}\}_{t=1}^{T}$ and $\{\tau_{it}\}_{t=1}^{T}$ can be changed by simply adding and subtracting an arbitrary constant without affecting the data. We impose $\tau_{c1} = 0$ to eliminate this ambiguity. Similarly, $\alpha_{\tau,it}\tau_{ct}$ and $\alpha_{\varepsilon,it}\varepsilon_{ct}$ are invariant to multiplying $\alpha_{\tau,it}$ or $\alpha_{\varepsilon,it}$ by a constant and dividing $\tau_{ct}, \sigma_{\Delta\tau,ct}, \varepsilon_{ct}, \sigma_{\varepsilon,ct}$ accordingly. Hence we fix $\sigma_{\Delta\tau,c1} = \sigma_{\varepsilon,c1} = 1$. See Del Negro and Otrok [2008] for further discussion on the normalizations.

Estimation Our model has data $w_t = \{w_{it}\}_{i=1}^n$, time-invariant parameters

$$\theta = \left(\{\theta_{c\ell}\}_{1\leq \ell \leq p}, \{\theta_{i\ell}\}_{1\leq \ell\leq q, 1\leq i\leq n}\right),$$
$$\gamma = \left(\{\gamma_{\alpha,m,i}\}_{m=\tau,\varepsilon,1\leq i\leq n}, \{\gamma_{\sigma,m,j}\}_{m=\Delta\tau,\varepsilon,\ j=c,1,\ldots,n}\right),$$

time-varying parameters

$$\lambda_t = \left(\{\alpha_{\tau,it}, \alpha_{\varepsilon,it}\}_{i=1}^n, \sigma_{\Delta\tau,ct}, \{\sigma_{\Delta\tau,it}\}_{i=1}^n, \sigma_{\varepsilon,ct}, \{\sigma_{\varepsilon,it}\}_{i=1}^n\right),$$

and latent variables

$$\xi_t = \left(\tau_{ct}, \{\tau_{it}\}_{i=1}^n, \varepsilon_{ct}, \{\varepsilon_{it}\}_{i=1}^n\right).$$

We use Bayesian methods to estimate the parameters and make inferences about the latent components. Specifically, we formulate a prior on (θ, γ) and we simulate a Markov chain $\{\theta^{(s)}, \gamma^{(s)}, \{\lambda_t^{(s)}\}, \{\xi_t^{(s)}\}\}_{s=1}^S$ that possesses as invariant distribution the joint posterior given the data $\{w_t\}$. We can then use the draws $\{\xi_t^{(s)}\}_{s=1}^S$ to form point and set estimates of the aggregate trend $\tilde{\tau}_t$, the sector-level trends $\{\tilde{\tau}_{it}\}_{i=1}^n$, and decomposition of trend changes over time in terms of common and sector-specific contributions. The priors and the Markov Chain Monte Carlo (MCMC) algorithm used for estimation and filtering are described in detail in Appendix B. Monte Carlo simulations showing that our estimation algorithm performs well in small samples are reported in Appendix C.

It is worth highlighting that our framework can accommodate missing wage observations in a straightforward manner. In the empirical analysis, we focus on a sample with no missing data that begins in 1997. The CPS allows us to estimate our model on a longer sample that begins in 1982 but has missing wage observations in the 80s and 90s. The results for the most recent period are robust to using either the longer or shorter sample, as we show in Appendix G.

Validation To validate our framework, we conduct a real-time out-of-sample forecast comparison against a random walk forecast computed on the wage data w_{it} at the aggregate and industry levels. As noted above, two features distinguish our approach: (i) the explicit treatment of temporal aggregation in the persistent component and (ii) the use of the cross-sectional dimension. The forecasting exercise we present here is designed to assess both. Since w_{it} as defined in Equation (1) is itself a 12-month growth rate of wages, the random walk forecast is comparable to the 12-month (or 4-quarter) forecasting rule that was found difficult to improve upon in the price inflation literature (see, e.g., Atkeson and Ohanian [2001], Stock and Watson [2008]).[12]

We construct a total of 300 datasets, each of them extending from January 1997 to T where T ranges from October 1998 to September 2023. For each dataset (indexed by T), we run our model and recover point estimates (the posterior median) of the path of the trend in wage growth at the aggregate $\hat{\tau}_t^T$ and at the sectoral level $\{\hat{\tau}_{it}^T\}_{i=1}^n$. We then define the following forecasts based on our model

$$f_{\text{TWIn}}^T(h) = \frac{\sum_{\ell=0}^{12-h} \hat{\tau}_{t-\ell}^T + h\hat{\tau}_t^T}{12}, \qquad f_{\text{TWIn},i}^T(h) = \frac{\sum_{\ell=0}^{12-h} \hat{\tau}_{i,t-\ell}^T + h\hat{\tau}_{it}^T}{12},$$

if the forecasting horizon $h < 12$ and $f_{\text{TWIn}}^T(h) = \hat{\tau}_t^T$, $f_{\text{TWIn},i}^T(h) = \hat{\tau}_{it}^T$ if $h \geq 12$. Thus, $f_{\text{TWIn}}^T(h)$ and $f_{\text{TWIn},i}^T(h)$ are point forecasts of 12-month wage growth at time $T + h$ given information available at T. We compare $f_{\text{TWIn}}^T(h)$ and $f_{\text{TWIn},i}^T(h)$ with random-walk forecasts, $f_{\text{RW}}^T(h) = w_T$ and $f_{\text{RW},i}^T(h) = w_{iT}$. Let $\text{RMSFE}_c(h)$ be the population root mean-square forecast errors (RMSFE) at horizon h for $c = \text{TWIn}, \text{RW}$ and let $\hat{R}_c(h) = \sqrt{\sum_{T=1}^N (w_{T+h} - f_c^T(h))^2 / N}$ be the sample RMSFE where $N = 300$ is the number of out-of-sample forecasts. We test

$$\text{RMSFE}_{\text{TWIn}} \geq \text{RMSFE}_{\text{RW}}$$

[12] We also considered as an alternative a 12-month moving average of w_{it}, denoted $\bar{w}_{it} = \sum_{\ell=1}^{12} w_{i,t+1-\ell}$, in order to compare our approach with a simple rule that smooths out the noise in w_{it} but this was dominated by the random-walk benchmark we report below.

by means of the one-sided test that rejects the null when the bootstrap estimate of the probability $P\left[\sqrt{N}\left(\hat{R}_{\mathrm{TWIn}}(h) - \hat{R}_{\mathrm{RW}}(h)\right) < 0\right]$ falls under the significance level α.[13] This is a high bar: we only reject the null when the evidence in favor of TWIn is sufficiently strong, whereas not rejecting the null does not rule out $\mathrm{RMSFE}_{\mathrm{TWIn}} = \mathrm{RMSFE}_{\mathrm{RW}}$.

Table 1 shows that $f^T_{\mathrm{TWIn}}(h)$ consistently outperforms $f^T_{\mathrm{RW}}(h)$ for a selection of short and long horizons h, as evidenced by the root mean square forecast error (RMSFE). The improvement is particularly clear for aggregate wage growth at the 3-month horizon, decreasing with h thereafter.

At the sector level, $f^T_{\mathrm{TWIn},i}$ is a substantially more accurate forecast than $f^T_{\mathrm{RW},i}$. The improvements occur for all sectors and horizons, but they are largest for the sectors that generally experience more variability in wage growth. Table 1 illustrates this for *construction & mining* and for *leisure & hospitality*, two sectors that received a lot of attention during the Great Financial Crisis and the COVID-19 Pandemic, and coincidentally the two sectors with the most volatile wage growth.

4 Empirical analysis

In the first part of this section, we present our main empirical results from the model estimated with our baseline industry partition. We then show that our model delivers a measure of wage inflation that is timelier than alternatives, as it strongly co-moves with indicators of labor market tightness. Finally, we show how our conclusions are broadly unaffected when re-estimating the model with alternative partitions of the data.

[13] This test is in the spirit of the Diebold and Mariano [1995] test except that (i) it is one-sided and (ii) we use a bootstrap procedure instead of a heteroskedasticity and autocorrelation consistent (HAC) estimate of the standard error of $\hat{S}_N$ together with standard normal critical values. Our implementation relies on the moving block bootstrap (Künsch [1989], Liu and Singh [1992]) with block size $B = \lceil N^{1/3} \rceil = 7$ following the discussion about rate-optimal block bootstrap in Lahiri [1999].

TABLE 1. Forecast comparison

	RMSFE			
	$h = 3$	$h = 6$	$h = 9$	$h = 12$
$f^T_{\text{TWIn}}(h)$	0.507	0.662	0.803	0.929
$f^T_{\text{RW}}(h)$	0.578	0.682	0.820	0.936
Difference	-0.071	-0.020	-0.017	-0.007
p-value	(0.00)	(0.15)	(0.21)	(0.16)
Construction and Mining				
$f^T_{\text{TWIn},i}(h)$	1.906	1.978	1.919	2.001
$f^T_{\text{RW},i}(h)$	2.583	2.680	2.456	2.451
Difference	-0.676	-0.702	-0.537	-0.450
p-value	(0.00)	(0.00)	(0.00)	(0.00)
Leisure and Hospitality				
$f^T_{\text{TWIn},i}(h)$	1.736	1.848	1.947	2.113
$f^T_{\text{RW},i}(h)$	2.248	2.198	2.280	2.450
Difference	-0.512	-0.350	-0.333	-0.337
p-value	(0.00)	(0.00)	(0.00)	(0.00)

NOTES. We report p-values for a one-sided test of the null hypothesis that the RMSFE of our model forecast is not lower than the RMSFE of random-walk forecasts. We compute the p-values by moving block bootstrap with blocks of size 7.

4.1 The dynamics of Trend Wage Inflation

We estimate the model breaking down wage inflation into seven broad industry groups. w_{it} is defined as in Equation (1) and thus is median nominal wage growth in industry i. Figure 2a shows our estimates of Trend Wage Inflation (TWIn, solid blue line) together with the realized 12-month wage growth from the CPS data (black line). A 68 percent posterior probability band is given by the blue shaded areas around TWIn.

We highlight two main takeaways. First, the underlying data exhibit a lot of monthly variation that is purged by our approach. Most of the high-frequency variation in nominal wage growth is ascribed to measurement error in our model (the $\tilde{\varepsilon}_{it}$ terms), as discussed in Section 3. Figure 2a makes clear that substantial movements in the persistent component

FIGURE 2. Trend Wage Inflation and its Decomposition

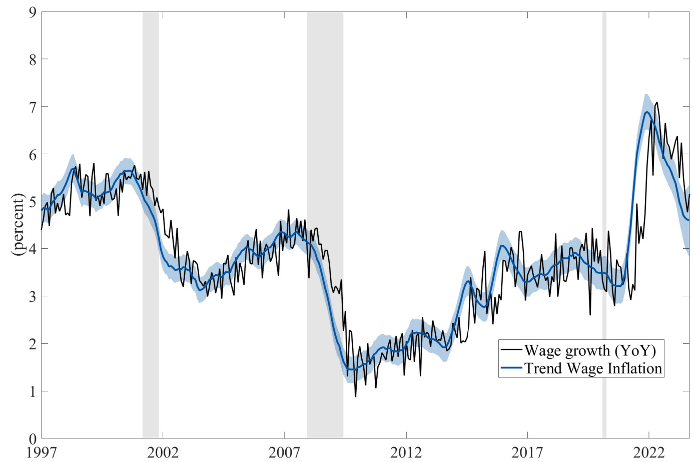

(a) Estimates of Trend Wage Inflation

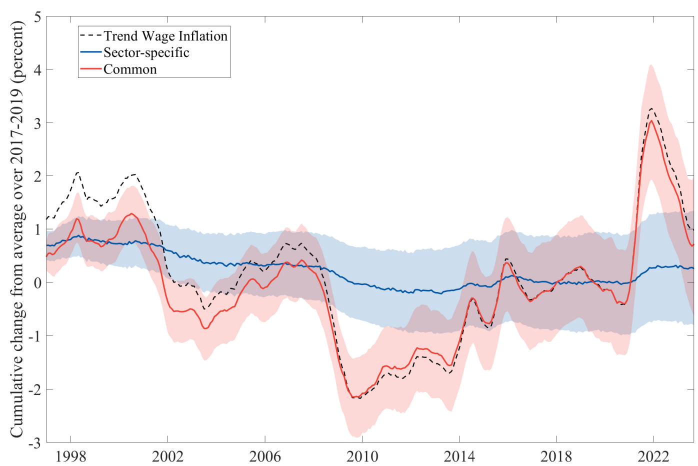

(b) Common and Sector-Specific Components of Trend Wage Inflation

NOTES. (a) Blue shaded areas denote 68 percent probability bands; grey shaded areas indicate recessions. (b) Shaded areas denote 68 percent probability bands. The dashed black line is nominal wage growth, the red line is the common component $\left(\sum_{i=1}^{n} s_{it}\alpha_{\tau,it}\right)\tau_{ct}$, while the blue line is the sector-specific component $\sum_{i=1}^{n} s_{it}\tau_{it}$. Each of these three series are plotted as a cumulative change from their average over the period 2017–2019, which is therefore centered at 0.

of wage inflation coincide with significant macroeconomic events, such as the Great Recession and the inflation surge episode starting in 2021.

As an example, we can analyze the wage inflation surge episode starting in 2021 through the lens of our model. TWIn ranged between 3.2 percent and 3.7 percent between

2016 and 2020, and then shot up in early 2021, nearly doubling over the course of the year. As such, the model assigns almost all of the nominal wage growth observed in the data to the persistent component.[14] Our TWIn estimates peak in December 2021 and have been slowly declining thereafter, although they were still above their pre-pandemic level in September 2023.

Second, we use our framework to further investigate whether fluctuations in TWIn are the result of sector-specific shocks. Recall from Equation (3) that TWIn can be written as the sum of a common component $\left(\sum_{i=1}^{n} s_{it}\alpha_{\tau,it}\right)\tau_{ct}$ and a sector-specific component $\sum_{i=1}^{n} s_{it}\tau_{it}$. Figure 2b depicts the evolution of each of these components over time.

We find that most of the variation in Trend Wage Inflation is explained by the common persistent component of nominal wage growth, both during the two NBER recessions and the inflation surge episode in the sample.[15] By contrast, the estimated sector-specific trend component of the model captures lower frequency movements. We find that it steadily falls over the period, which implies an increase in the cross-sector correlation of nominal wage growth. Over the sample period, the sector-specific component especially declines for Education and Health, and Finance and Business services. We tentatively conjecture that this pattern could be related to college-educated workers, who tend to be over-represented in these industries, since the college wage premium was still on an increasing trajectory in the period leading up to the Great Recession [Acemoglu and Autor, 2011]. The sector-specific component of TWIn increases following the pandemic (although the probability bands are large), but it remains below its level of the early 2000s.

Our estimates imply that variations in the industry-specific components are second-order for aggregate fluctuations in Trend Wage Inflation. However, they do not imply that changes in wage inflation are the same in all industries since the loadings α_{it} are

[14] As discussed in Section 3, the trend extracted by the model is expressed in terms of annualized monthly wage growth, which explains why it leads the actual year-over-year wage growth series in the chart.

[15] This finding is consistent with the analysis of sectoral mismatch during the Great Recession in Şahin et al. [2014]. They find limited evidence of sectoral mismatch between job seekers and job openings. While wages are not analyzed in their study, this result is in line with sectoral wage inflation being driven by a common trend.

specific to each industry, and their "sensitivity" to the common factor can therefore differ.[16] As a result, we still find that the common factor is preeminent for industries that can be expected to be specifically impacted at some point in our sample, such as construction during the Great Recession and leisure and hospitality following the COVID-19 pandemic (see Appendix D.1). One exception is Public Administration, which features quantitatively important sector–specific dynamics. This is consistent with micro-level evidence suggesting that wages in the public sector take longer to adjust (Barattieri et al. [2014]).

To sum up, our framework assigns most of the business cycle variation in nominal wage inflation to Trend Wage Inflation. In turn, TWIn is mostly driven by the component that is common across industries.

4.2 TWIn as a timely indicator of wage pressures

Our filtering approach to time aggregation delivers a measure of wage inflation that is timelier than alternatives. In this section, we show this in two ways.

First, in Figure 3, we plot TWIn, the AWGT, and the growth rate in ECI together with labor market tightness, for the three most significant macroeconomic episodes of the last 25 years.[17] Our measure of Trend Wage Inflation always leads alternative measures of wage growth: importantly, it is better aligned to labor market tightness, moving almost contemporaneously.[18] As described in Section 3, our estimates are derived purely from

[16] Loadings can also change over time. In Appendix D.2, we show that these movements happen at relatively low frequency. In unreported results, we also find that the conclusions of our paper are broadly unaffected if we keep the loadings constant over time. Appendix D.2 also shows that stochastic volatilities change over time, albeit slowly, underscoring the importance of allowing time-varying parameters in our model.

[17] TWIn co-moves with labor market tightness even outside of these episodes. For instance, it increased since 2014, together with tightness, while other wage growth measures stalled. A similar behavior is also displayed by the indicator of common wage inflation proposed by Ahn et al. [2024].

[18] In standard job search models, wages and labor market tightness typically move together [see Pissarides, 2000]. We note that the amplitude of TWIn tends to be greater than that of labor market tightness, especially for the 2007 and 2021 episodes. We do not see this as a failure of our indicator since, to the best of our knowledge, there is no theoretical reasons for the rate of changes of wage growth and labor market tightness to be the same.

wage data, so the patterns shown in Figure 3 are the direct result of taking into account the time aggregation implied by 12-month changes in our definition of wage inflation (1).[19] Because the underlying wage data are released on a timely basis, this property of our measure also applies in close to real time.[20]

Second, we find that changes in TWIn (and in its common component) are significantly correlated with changes in labor market conditions and price inflation measures, confirming that TWIn can be seen as a reliable real-time indicator of wage pressures. Importantly, these correlations are stronger not only relative to raw CPS data, but also relative to other popular measures of aggregate wage growth, as we show in Table 2. In line with the previous discussion, the vacancy to labor force ratio strongly and significantly co-moves with TWIn; that correlation is insignificant for ECI, absent for the AWGT, and even of the opposite sign for AHE. A similar pattern is observed for the unemployment rate and the negative unemployment rate gap, although differences are less stark – with the exception of AHE. In unreported results, we also find that TWIn is also strongly and significantly correlated with the negative unemployment rate gap when we estimate wage Phillips curves as in Ahn et al. [2024]. Changes in TWIn are positively and significantly correlated with changes in inflation, consistent with the idea that wage pressures may be associated with price pressures. The ECI has a similar pattern but a weaker association; AWGT is uncorrelated with inflation whereas AHE has, again, the opposite sign.

4.3 Heterogeneity and aggregate dynamics of Trend Wage Inflation

In this section, we investigate the robustness of our finding that aggregate wage inflation is driven by a common component. We consider a series of alternative partitions of the data that represent likely sources of heterogeneity in nominal wage growth.

[19] In unreported results, we repeat this analysis without accounting for time-aggregation and find that the timeliness properties of our indicator are substantially degraded.

[20] The CPS data underlying our estimates are released by mid-month for the previous month. By comparison, JOLTS job openings, which are required to derive labor market tightness, are typically released with a month's lag.

FIGURE 3. Labor Market Tightness and Wage Inflation

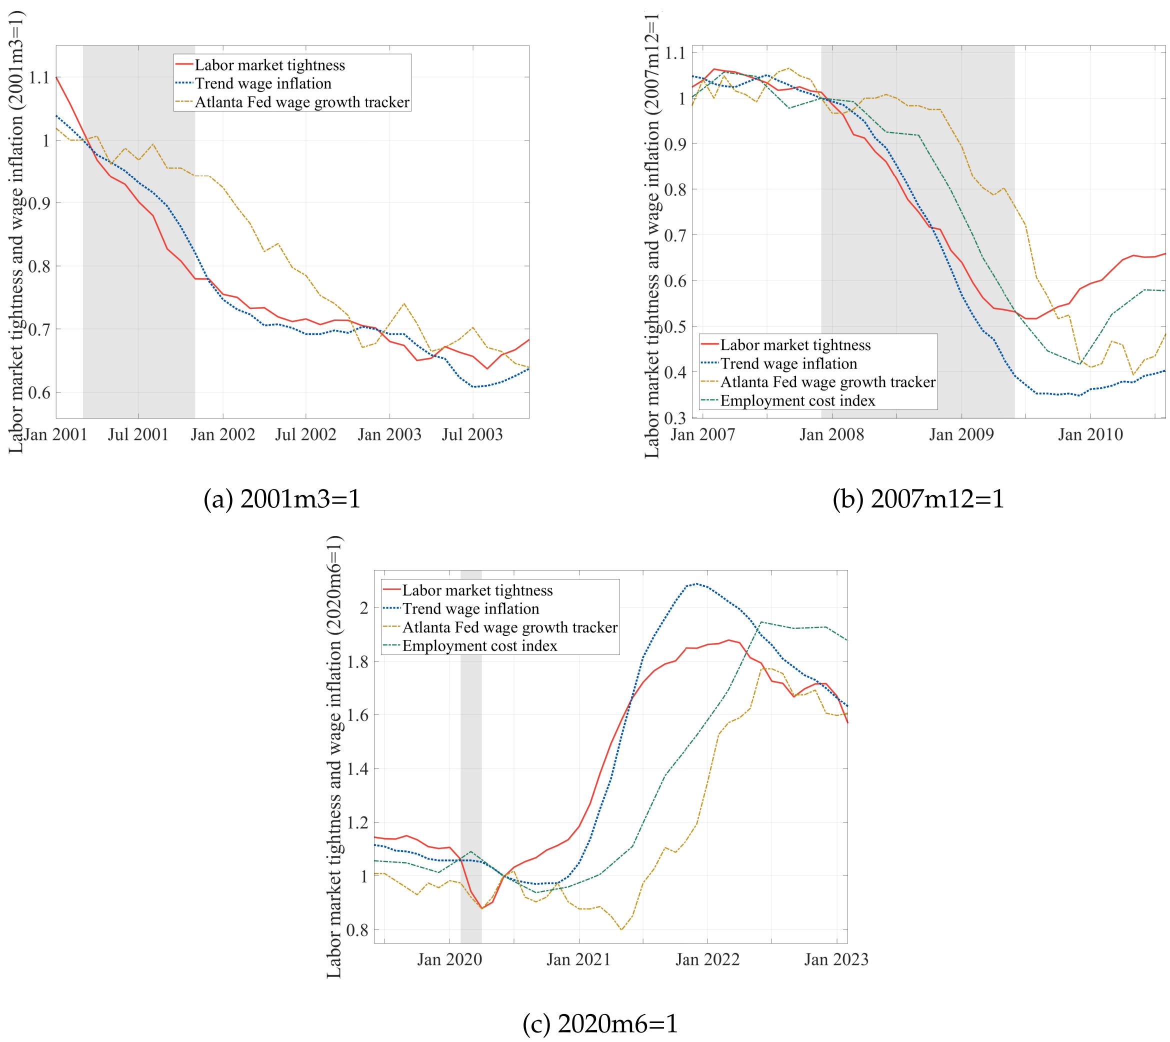

NOTES. The red solid line represents labor market tightness measured by the ratio of vacancies to the sum of employed and unemployed workers. We apply a symmetric one lag-one lead moving average to smooth it. Dotted lines are Trend Wage Inflation (blue), the Atlanta Fed Wage Growth Tracker (yellow) and the year-on-year change in the employment cost index (green). The Employment Cost Index is only available starting in 2001q1, so the four-quarter change starts can is only available from 2002q1.

Our measure of wage inflation comes from the aggregation of changes in nominal wages at the worker level, which do not occur frequently. Many explanations for these nominal wage rigidities have been put forward in the literature. Examples include employers insuring workers against firm-level shocks, fairness considerations related to efficiency, and institutionalized wage setting, such as employers setting wages at the na-

TABLE 2. Correlations with labor market and price inflation time series

	Vacancy to labor force ratio	Unemployment rate	(−) UR gap	Core PCE inflation	Core services ex housing PCE inflation
Monthly wage inflation measures					
TWIn	0.331***	−0.094*	0.191*	0.300***	0.242***
TWIn (common)	0.320***	−0.077	0.183*	0.310***	0.245***
AHE	−0.174***	0.588***	−0.550***	−0.227***	−0.226***
Atlanta Fed wage tracker	−0.003	−0.115**	0.158	0.088	0.028
Quarterly wage inflation measures					
TWIn	0.718***	−0.295***	0.221**	0.427***	0.334***
TWIn (common)	0.717***	−0.285***	0.210**	0.436***	0.334***
ECI	0.159	−0.158	0.130	0.221**	0.209**

NOTES. TWIn, Atlanta Fed wage tracker, AHE, the unemployment rate and price inflation measures are monthly series over the period 1997m1–2023m9. The vacancy ratio is over the period 2000m12–2023m9. Results are qualitatively similar prior to 2020m3. Vacancies are seasonally adjusted job openings from the Job Openings and Labor Turnover Survey (JOLTS) of the Bureau of Labor Statistics. The negative UR gap is the difference between the short-term natural rate of unemployment from the Congressional Budget Office (CBO) and the unemployment rate, constructed as in Ahn et al. [2024]. AHE is the 12-month percent change in average hourly earnings of production and non-supervisory employees on private nonfarm payrolls, from the Current Employment Statistics of the Bureau of Labor Statistics. The Atlanta Fed wage tracker (Atlanta Fed [2023]) is the unweighted 3-month moving average of median 12-month wage growth. Core PCE inflation comes from the Bureau of Economic Analysis and excludes energy and food. ECI is a quarterly measure of the 12-month percent change in the Employment Cost Index measured by the Bureau of Labor Statistics. The time period is 1997Q1–2023Q2. When computing correlations at quarterly frequency, the price inflation measures and TWIn are 12-month changes using the third month of the quarter as ECI. All correlations are for variables in first differences. *, **, and *** denote statistical significance at the 10%, 5%, and 1%, respectively.

tional level or minimum wage institutions.[21] All of these explanations are likely to result in different degrees of nominal wage stickiness for separate groups of workers. As a concrete example, national wage setting by employers or minimum wage regulations are more likely to apply to low-skill workers.[22]

To assess whether this micro-level heterogeneity affects our earlier discussion on *aggregate* wage inflation, we re-estimate our model for a series of partitions that we derive from

[21]On insurance motives, see for instance a seminal paper by Barro [1977] and Azariadis and Stiglitz [1983] for a review. Early papers on the theory of efficiency wages include Stiglitz [1976] and Akerlof and Yellen [1990].

[22]See Avouyi-Dovi et al. [2013] and Minton and Wheaton [2022] for recent work on minimum wages and downward wage rigidity. Hazell et al. [2022] show how national wage setting increases wage rigidity.

the CPS. Besides the industry partition discussed in Section 4.1, we consider partitions based on occupation, education level, region, wage quartile, age, gender, and race.[23] For each group, we obtain estimates of Trend Wage Inflation and its common and idiosyncratic components given in Equation (3).

We summarize our estimation results in Table 3, focusing on the three largest changes in TWIn in the sample: the 2001 Recession (2001m3-2001m11), the Great Recession (2007m12-2009m6), and the post-pandemic wage inflation surge episode (2020m6-2022m2).[24] We report the fraction of the cumulative change in TWIn accounted for by its common component over the duration of each episode. We find that the common component remains an important determinant of TWIn across the various partitions, but to a lesser extent during the 2001 recession (36 percent to 91 percent) and the Great Recession (50 percent to 90 percent). Industries are measured to have a larger contribution of the common component of wage inflation than any other cut of the data across the three episodes considered (greater than 90 percent overall), suggesting that most of the industry-specific trend movements reflect differential loadings (the α's in our model). Our estimates suggest that the common component of wage inflation accounts for a lower share of the evolution of wages during these episodes when i is defined using demographic variables. This finding is consistent with the literature on the heterogeneous impact of recessions across demographic groups. For instance, prior work has shown that female labor supply tends to be less cyclical than that of men during US recessions [Doepke and Tertilt, 2016]—a finding sometimes labeled "man-cession." Our indicator suggests that there is similar heterogeneity in wage inflation data.

Our result on the importance of the common component of the trend is not a limitation of the econometric framework because some sectors are estimated to have quantitatively

[23] We follow the definitions used in the Atlanta Fed Wage Growth Tracker. See Appendix A for details. In unreported results, we also find that our results are robust to interactions of groups (e.g.: age and occupation).

[24] The recession episodes follow NBER dating – with start date one month before the onset of the recession. The inflationary episode reflects trough-to-peak dates of core PCE inflation. The results are robust to considering slightly different start and end dates.

TABLE 3. Contribution of common trend to Trend Wage Inflation during selected episodes

	2001 Recession	Great Recession	2021 Inflation Surge
Industry	0.91	0.90	0.92
	(0.65, 1.17)	(0.76, 1.02)	(0.81, 1.03)
Occupation	0.58	0.71	0.86
	(0.25, 0.87)	(0.49, 0.92)	(0.70, 1.02)
Education	0.62	0.74	0.83
	(0.31, 0.94)	(0.52, 0.95)	(0.67, 1.00)
Region	0.83	0.86	0.94
	(0.51, 1.09)	(0.73, 0.98)	(0.83, 1.04)
Wage quartile	0.70	0.75	0.89
	(0.34, 1.05)	(0.56, 0.93)	(0.73, 1.04)
Age	0.61	0.74	0.87
	(0.19, 0.98)	(0.54, 0.95)	(0.68, 1.05)
Gender	0.43	0.54	0.82
	(0.08, 0.80)	(0.28, 0.81)	(0.60, 1.03)
Race	0.36	0.50	0.84
	(0.03, 0.77)	(0.20, 0.78)	(0.62, 1.03)

NOTES. Each row refers to a different model, using the cross–sectional variable reported, defined following the Atlanta Fed Wage Growth tracker definition and detailed in Appendix A. The three episodes in the columns refer to the following periods: 2001m3-2001m11, 2007m12-2009m6, and 2020m6-2022m2. For each partition and episode, we compute the change in the common component of Twin between the start and end of the period, as well as the change in the overall TWIn over the same horizon, and report the median of this ratio. We report a 68 percent posterior probability interval in parentheses.

important sector-specific trends. One example we have already mentioned is Public Administration in the industry-level model. As an additional check, Appendix C reports Monte Carlo simulations showing that our method is not biased toward attributing an excessive role to the common component.

Notably, we find that the common component of wage inflation accounts for most of the increase in TWIn across the various partitions during the 2021 inflation surge episode. The contribution of the common component is greater than 82 percent across subgroups during that episode, typically higher than during the the two recessions in the sample.[25]

[25] In Appendix D.3 we show that the estimated change in TWIn in 2021 is very robust across specifications.

This finding hints at the existence of a form of asymmetry between recessionary and inflationary episodes, with group-specific forces being more important during economic downturns. To investigate whether this result is specific to the 2021 inflation surge, the only high-inflation episode in the CPS sample, we also estimate our model using average hourly earnings data from the Current Employment Statistics (CES). Despite the composition issues discussed in Section 2, the CES has the advantage to start in 1964.[26] As shown in Figure 4, which repeats the analysis of Figure 2b in this longer sample, the common persistent component is also the main driver of wage inflation during the inflation episodes of the 1970s. For instance, it accounts for 94 percent of the 1.8 percentage points increase in the trend over the period 1973-74. By contrast, industry-specific variation accounts for 30 percent of the 0.7 percentage point drop in the overall trend during the Great Recession.

FIGURE 4. Estimates using CES data

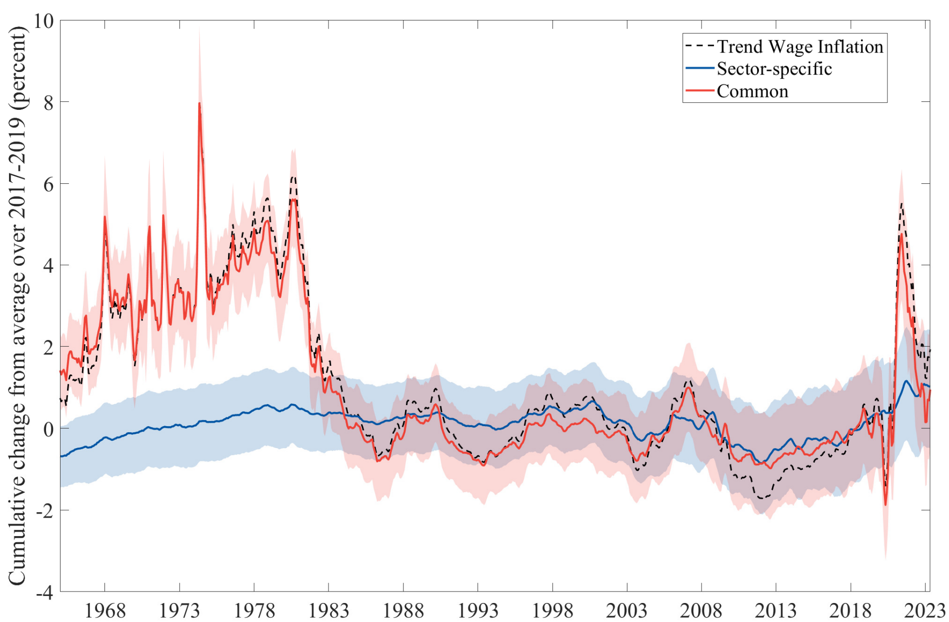

To sum up, no specific subgroup, in the eight partitions that we consider, seems

[26]See Appendix F for additional details.

to play a disproportionate role in driving the overall trend in nominal wage growth, especially during inflationary periods. A useful case study is the model estimated with wage quartiles. Recent work by Autor et al. [2023] suggests that the COVID-19 period has seen substantial wage compression, with nominal wage growth at the bottom of the distribution being relatively stronger than in the upper echelons. While our estimated model suggests that nominal wage growth in the bottom wage quartile is initially driven by group-specific forces, this has a relatively minor role for the evolution of wage growth in the aggregate during the post-pandemic inflation surge.

We want to reiterate that our findings do not imply that there is no heterogeneity in wage stickiness at the micro level, including during inflationary episodes. Instead, our framework suggests that group-specific movements in nominal wage growth mostly reflect differential sensitivity to an aggregate factor, rather than sector-specific trends. By estimating this common factor, our model effectively isolates the key driver of aggregate nominal wage growth.

5 Conclusion

We propose a new framework to measure the latent trend component of monthly wage inflation by combining worker–level data with time series smoothing methods that account for the temporal aggregation in the CPS. The resulting measure of Trend Wage Inflation (TWIn) allows us to establish four distinct empirical facts.

First, TWIn fluctuates substantially during episodes of macroeconomic relevance (such as recessions and inflationary spikes). Second, most of its variation is driven by a common component indicating a strong degree of cross-sectional correlation across workers. Taken together, these two facts imply that most of the business cycle fluctuations in wage inflation are driven by a persistent component that is common across workers. Third, changes in TWIn are strongly contemporaneously correlated with changes in measures of labor market tightness and price inflation, more so than existing indicators of nominal wage growth. This makes our measure a reliable real-time indicator of wage pressures. And fourth, observable heterogeneity in worker and job characteristics, although important for the level and sensitivity to the common component, is second order for the aggregate trend in wage growth.

These facts provide an empirical basis against which to test models of the determination of wage inflation and labor market tightness that incorporate aggregate shocks and heterogeneity. One avenue for future research is to perform such tests. Another is to replicate our empirical analysis in other countries, particularly in those with different labor market institutions and wage-setting practices. Finally, since wage inflation is related to changes in the marginal cost of labor, yet another avenue is to develop a flexible empirical joint model of the persistent components of price and wage inflation.

References

Acemoglu, D. and Autor, D. (2011). Chapter 12-skills, tasks and technologies: Implications for employment and earnings (d. card & o. ashenfelter, eds.). *Elsevier. https://doi. org/10.1016/S0169-7218 (11)*, pages 02410–5.

Ahn, H. J., Chen, H., and Kister, M. (2024). A new indicator of common wage inflation. *Journal of Money, Credit and Banking*, Forthcoming.

Akerlof, G. A. and Yellen, J. L. (1990). The fair wage-effort hypothesis and unemployment. *The Quarterly Journal of Economics*, 105(2):255–283.

Atkeson, A. and Ohanian, L. E. (2001). Are Phillips curves useful for forecasting inflation? *Quarterly Review*, 25(Win):2–11.

Atlanta Fed (2023). Atlanta fed wage growth tracker. `https://www.atlantafed.org/chcs/wage-growth-tracker`. Accessed: 2023-04-30.

Autor, D., Dube, A., and McGrew, A. (2023). The unexpected compression: Competition at work in the low wage labor market. Technical report, National Bureau of Economic Research.

Avouyi-Dovi, S., Fougère, D., and Gautier, E. (2013). Wage rigidity, collective bargaining, and the minimum wage: evidence from french agreement data. *Review of Economics and Statistics*, 95(4):1337–1351.

Azariadis, C. and Stiglitz, J. E. (1983). Implicit contracts and fixed price equilibria. *The Quarterly Journal of Economics*, pages 2–22.

Barattieri, A., Basu, S., and Gottschalk, P. (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics*, 6(1):70–101.

Barro, R. J. (1977). Long-term contracting, sticky prices, and monetary policy. *Journal of Monetary Economics*, 3(3):305–316.

Bils, M. J. (1985). Real wages over the business cycle: evidence from panel data. *Journal of Political economy*, 93(4):666–689.

Daly, M. C., Hobijn, B., and Wiles, T. S. (2011). Aggregate real wages: macro fluctuations and micro drivers. Technical report, Tinbergen Institute Discussion Paper.

Del Negro, M. and Otrok, C. (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. Technical report.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.

Doepke, M. and Tertilt, M. (2016). Families in macroeconomics. In *Handbook of macroeconomics*, volume 2, pages 1789–1891. Elsevier.

Gertler, M., Huckfeldt, C., and Trigari, A. (2020). Unemployment fluctuations, match quality, and the wage cyclicality of new hires. *The Review of Economic Studies*, 87(4):1876–1914.

Ghysels, E. (2012). Mixed data sampling regression models. In Elliott, G. and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2, pages 347–386. Elsevier.

Ghysels, E., Kvedaras, V., and Zemlys-Balevičius, V. (2020). Mixed data sampling (midas) regression models. In Sengupta, D. K. and Rao, C., editors, *Handbook of Statistics*, volume 42, pages 117–153. Elsevier.

Grigsby, J., Hurst, E., and Yildirmaz, A. (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111(2):428–71.

Hazell, J., Patterson, C., Sarsons, H., and Taska, B. (2022). National wage setting. Technical report, National Bureau of Economic Research.

Hazell, J. and Taska, B. (2020). Downward rigidity in the wage for new hires. *Available at SSRN 3728939*.

Huh, C. G. and Trehan, B. (1995). Modeling the time-series behavior of the aggregate wage rate. *Economic Review-Federal Reserve Bank of San Francisco*, (1):3.

Kiley, M. T. (2023). The role of wages in trend inflation: Back to the 1980s?

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*, 17(3):1217–1241.

Lahiri, S. N. (1999). Theoretical comparisons of block bootstrap methods. *The Annals of Statistics*, 27(1):386–404.

Lewis, D. J., Mertens, K., Stock, J. H., and Trivedi, M. (2022). Measuring real activity using a weekly economic index. *Journal of Applied Econometrics*, 37(4):667–687.

Liu, R. Y. and Singh, K. (1992). Moving blocks jackknife and bootstrap capture weak dependence. In Lepage, R. and Billard, L., editors, *Exploring the Limits of Bootstrap*, pages 224–248. John Wiley.

Minton, R. and Wheaton, B. (2022). Minimum wages and the rigid-wage channel of monetary policy1.

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.

Şahin, A., Song, J., Topa, G., and Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, 104(11):3529–3564.

Stiglitz, J. E. (1976). The efficiency wage hypothesis, surplus labour, and the distribution of income in ldcs. *Oxford economic papers*, 28(2):185–207.

Stock, J. H. and Watson, M. W. (2008). Phillips curve inflation forecasts. Working Paper 14322, National Bureau of Economic Research.

Stock, J. H. and Watson, M. W. (2016a). Core inflation and trend inflation. *Review of Economics and Statistics*, 98(4):770–784.

Stock, J. H. and Watson, M. W. (2016b). Core inflation and trend inflation. *Review of Economics and Statistics*, 98(4):770–784.

Watson, M. W. and Engle, R. F. (1981). A one-factor multivariate time series model of metropolitan wage rates. *Journal of the American Statistical Association*, 76(376):774–781.

Supplemental Appendix to

"A Measure of Trend Wage Inflation"

by Martín Almuzara[†], Richard Audoly[‡], and Davide Melcangi[§]

A Definition of CPS partitions

We partition workers in the CPS using the same definitions as the Atlanta Fed WGT [Atlanta Fed, 2023].

Industries (7 groups) Construction and Mining, Education and Health, Finance and Business Services, Leisure and Hospitality, Manufacturing, Public Administration, and Trade and Transportation.

Occupations (3 groups) high–skill (Managers, Professionals, Technicians), middle–skill (Office and Administration, Operators, Production, Sales), and low–skill (Food Preparation and Serving, Cleaning, individual Care Services, Protective Services).

Race (2 groups) White and Nonwhite.

Education (3 groups) High school or less, Associates degree, and Bachelor degree or higher.

Age (3 groups) 16–24 years old, 25–54, and 55+.

[†]Federal Reserve Bank of New York. Email: martin.almuzara@ny.frb.org. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

[‡]Federal Reserve Bank of New York. Email: richard.audoly@ny.frb.org.

[§]Federal Reserve Bank of New York. Email: davide.melcangi@ny.frb.org.

Gender (2 groups) Male and Female.

Wage quartiles (4 groups) The quartiles are based on the average between workers' current hourly wage and their wage 12 months prior (when their wages are last recorded).

Region (9 groups) The nine Census Divisions: New England (Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, Vermont), Mid-Atlantic (New Jersey, New York, Pennsylvania), East North Central (Illinois, Indiana, Michigan, Ohio, Wisconsin), West North Central (Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota), South Atlantic (Delaware, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, District of Columbia, West Virginia), East South Central (Alabama, Kentucky, Mississippi, Tennessee), West South Central (Arkansas, Louisiana, Oklahoma, Texas), Mountain (Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, Wyoming), Pacific (Alaska, California, Hawaii, Oregon, Washington).

B Details of model and estimation approach

Recall our notation for the data $w_t = \{w_{it}\}_{i=1}^n$, the time-invariant parameters

$$\begin{aligned}\theta &= \left(\{\theta_{c\ell}\}_{1\leq\ell\leq p}, \{\theta_{i\ell}\}_{1\leq\ell\leq q, 1\leq i\leq n}\right),\\ \gamma &= \left(\{\gamma_{\alpha,m,i}\}_{m=\tau,\varepsilon,1\leq i\leq n}, \{\gamma_{\sigma,m,j}\}_{m=\Delta\tau,\varepsilon,\ j=c,1,\ldots,n}\right),\end{aligned}$$

the time-varying parameters

$$\lambda_t = \left(\{\alpha_{\tau,it}, \alpha_{\varepsilon,it}\}_{i=1}^n, \sigma_{\Delta\tau,ct}, \{\sigma_{\Delta\tau,it}\}_{i=1}^n, \sigma_{\varepsilon,ct}, \{\sigma_{\varepsilon,it}\}_{i=1}^n\right),$$

and the latent components

$$\xi_t = (\tau_{ct}, \{\tau_{it}\}_{i=1}^n, \varepsilon_{ct}, \{\varepsilon_{it}\}_{i=1}^n).$$

To conduct Bayesian inference, we begin by formulating a prior on (θ, γ).

Choice of priors The MA coefficients θ are prior independent of each other with $\theta_{j\ell} \sim N(0, v_\ell^2)$ for $j = c, 1, \ldots, n$. That is, we shrink the model towards one with white noise transitory errors and the strength of the shrinkage is determined by the choice of v_ℓ. In our baseline model we set $p = 0$ and $q = 3$, and we put higher penalties on the more distant lags as in the Minnesota prior of Doan, Litterman, and Sims [1983]. We achieve that by setting $v_\ell = 1/(10\ell^2)$.

The standard deviations γ control the amount of time-variation in loadings and volatilities. Unless they are small, the model may be excessively flexible causing overfitting. Our approach is to put a reasonably tight prior centered around small values to shrink the model towards no time-variation in the parameters. Specifically we use independent inverse gamma priors of the form $\gamma^2_{k,m,j} \sim \Gamma^{-1}(d_k/2, 2/(d_k s_k^2))$ for $k = \alpha, \sigma$. The location parameters are set to $s_\alpha^2 = 0.0001$ and $s_\sigma^2 = 0.001$, and the degree-of-freedom hyperparameters are set to $d_\alpha = d_\sigma = 60$.

Estimation and filtering Inference about parameters and latent variables is implemented via Gibbs sampling. This is a type of Markov Chain Monte Carlo (MCMC) algorithm suitable to approximate the joint posterior distribution of parameters and latent variables by simulation in state-space models.

The Gibbs sampler constructs a Markov Chain $\left\{\theta^{(s)}, \gamma^{(s)}, \{\lambda_t^{(s)}\}, \{\xi_t^{(s)}\}\right\}_{s=1}^{S}$ having as invariant distribution the posterior

$$P\left(\theta, \gamma, \{\lambda_t\}, \{\xi_t\} \middle| \{w_t\}\right).$$

This allows us to estimate the posterior of our objects of interest, e.g.

$$P\left(\{\tilde{\tau}_t, \{\tilde{\tau}_{it}\}_{i=1}^n\} \middle| \{w_t\}\right),$$

using the draws $\left\{\theta^{(s)}, \gamma^{(s)}, \{\lambda_t^{(s)}\}, \{\xi_t^{(s)}\}\right\}_{s=1}^{S}$ to form $\left\{\{\tilde{\tau}_t^{(s)}, \{\tilde{\tau}_{it}^{(s)}\}_{i=1}^{n}\}\right\}_{s=1}^{S}$ and taking the simulation frequencies of the objects as estimates of posterior probabilities. If the Markov chain converges (in a suitable sense) and S is large, the approximation error will be small.

One advantage of the Bayesian approach is that posterior calculations already integrate both the sampling uncertainty from parameter estimation and the signal-extraction uncertainty about the latent components. When reporting the path over time of a latent time series in our empirical analysis, we use credible intervals with fixed credibility level pointwise in t.[1]

An alternative would be to estimate (θ, γ) by maximum likelihood. It is straightforward to modify our MCMC algorithm to approximate the maximum likelihood estimator by stochastic EM — simply replace the posterior updates of θ and γ by the solutions to the corresponding complete-data score equations. However, inferences about latent variables (and, in particular, about our objects of interest) that arise from that procedure would not necessarily account for the estimation uncertainty in (θ, γ).

Gibbs sampling Our algorithm follows Stock and Watson [2016a] who build on the method proposed by Del Negro and Otrok [2008] to estimate dynamic factor models with time-varying loadings and volatilities.

Relative to Del Negro and Otrok [2008], Stock and Watson [2016a] incorporate outliers in the transitory shocks. Compared to Stock and Watson [2016a], we allow for temporal aggregation in the persistent components and for MA dynamics in the transitory components. For simplicity, we discuss estimation of a model without outliers.[2] We find only a negligible role for them in the data we analyze.

The Gibbs sampler exploits the fact that with a careful grouping of parameters and latent variables, the conditional distributions of each block given the rest can be simulated

[1] It is conceptually straightforward and computationally feasible to compute pathwise credible regions along the lines of, e.g., Inoue and Kilian [2016].

[2] The outliers in Stock and Watson [2016a] are introduced by assuming $\eta_{\varepsilon,jt} = s_{jt} \times \tilde{\eta}_{\varepsilon,jt}$ for $j = c, 1, \ldots, n$ where $s_{jt} = 1$ with probability p_j and $s_{jt} \sim U(1, 10)$ with probability $1 - p_j$ while $\tilde{\eta}_{\varepsilon,jt} \sim N(0, 1)$.

by well-known algorithms. In our model, there are three big blocks with many sub-blocks, namely:

(A) $\boldsymbol{P(\{\xi_t\}|\{\lambda_t\}, \theta, \gamma, \{w_t\})}$. Conditional on time-varying parameters $\{\lambda_t\}$ and the MA coefficients θ, the data w_t and the latent variables ξ_t are related by a linear state-space model with time-varying matrices. We apply the simulation smoother algorithm proposed by Durbin and Koopman [2002] to efficiently sample $\{\xi_t\}$.

To accommodate the MA dynamics of the common and sector-specific transitory errors, we include $\varepsilon_{ct}, \varepsilon_{c,t-1}, \ldots, \varepsilon_{c,t-p+1}, \{\varepsilon_{it}, \varepsilon_{i,t-1}, \ldots, \varepsilon_{i,t-q+1}\}_{i=1}^{n}$ as additional state variables.

(B) $\boldsymbol{P(\{\lambda_t\}|\{\xi_t\}, \theta, \gamma, \{w_t\})}$. This can be further partitioned into the following blocks:

(i) $\boldsymbol{P\left(\left\{\{\alpha_{\tau,it}, \alpha_{\varepsilon,it}\}_{i=1}^{n}\right\}\middle|\left\{\{w_{it}, \tau_{it}\}_{i=1}^{n}\right\}, \{\tau_{ct}, \varepsilon_{ct}\}, \left\{\{\sigma_{\varepsilon,it}\}_{i=1}^{n}\right\}, \{\gamma_{\alpha,\tau,i}, \gamma_{\alpha,\varepsilon,i}\}_{i=1}^{n}\right)}$. It is the result of a multivariate regression with time-varying coefficients and MA error terms. It can be dealt with using linear state-space techniques. Thus, we apply the simulation smoothing algorithm of Durbin and Koopman [2002] to the corresponding state-space representation in order to sample $\{\{\alpha_{\tau,it}, \alpha_{\varepsilon,it}\}_{i=1}^{N}\}$.

(ii) $\boldsymbol{P\left(\{\sigma_{m,jt}\}\middle|\{m_{jt}\}, \gamma_{\sigma,m,j}\right)}$ for $m = \Delta\tau, \varepsilon$ and $j = c, 1, \ldots, n$. Given $\gamma_{\sigma,m,j}$ and m_{jt}, $\sigma_{m,jt}$ follows a stochastic-volatility model with observation equation $\ln m_{jt}^2 = \ln \sigma_{m,jt}^2 + \ln \eta_{m,jt}^2$ and transition equation $\Delta \ln \sigma_{m,jt}^2 = \gamma_{\sigma,m,j} \nu_{\sigma,m,jt}$. We then use the algorithm proposed in Kim, Shephard, and Chib [1998] and Omori, Chib, Shephard, and Nakajima [2007] that consists of approximating the log-χ_1^2 distribution of $\ln \eta_{\sigma,m,jt}^2$ with a 10-component normal mixture and applying linear state-space techniques to that approximation.

(C) $\boldsymbol{P(\theta, \gamma|\{\xi_t\}, \{\lambda_t\}, \{w_t\})}$. This can also be partitioned into subblocks:

(i) $\boldsymbol{P\left(\gamma_{\alpha,m,i}\middle|\{\Delta\alpha_{m,it}\}\right)}$ for $m = \tau, \varepsilon$. We draw the reciprocal of the square root of a gamma random variable with $d_\alpha + T$ degrees of freedom and mean $(d_\alpha s_\alpha^2 +$

$\sum_{t=1}^{T} \Delta\alpha_{m,it})/(d_\alpha + T)$ for $m = \tau, \varepsilon$.

(ii) $\boldsymbol{P\left(\gamma_{\sigma,m,j} \middle| \{\Delta \ln \sigma^2_{m,jt}\}\right)}$ for $m = \Delta\tau, \varepsilon$ and $j = c, 1, \ldots, n$. We draw the reciprocal of the square root of a gamma random variable with $d_\sigma + T$ degrees of freedom and mean $(d_\sigma s_\sigma^2 + \sum_{t=1}^{T} \Delta \ln \sigma^2_{m,jt})/(d_\sigma + T)$ for $m = \tau, \varepsilon$ and $j = c, 1, \ldots, n$.

(iii) $\boldsymbol{P\left(\theta_j \middle| \{\varepsilon_{jt}\}, \{\sigma_{\varepsilon,jt}\}\right)}$ for $j = c, 1, \ldots, n$ where $\theta_c = (\theta_{c1}, \ldots, \theta_{cp})'$ and $\theta_i = (\theta_{i1}, \ldots, \theta_{iq})'$ for $i = 1, \ldots, n$. This problem can be treated separately for each j. We do the derivation for $j = i = 1, \ldots, n$ (the case $j = c$ is identical except that p should take the place of q). Define $v_{it} = \sigma_{\varepsilon,it}\eta_{\varepsilon,it}$. Conditioning on q initial observations $\varepsilon_{i0}, \ldots, \varepsilon_{iq-1}$ we obtain the likelihood term

$$
\begin{aligned}
P\left(\{\varepsilon_{it}\}_{t=1}^{T} \middle| \{\theta_{i\ell}\}_{\ell=1}^{q}, \{\varepsilon_{i1-\ell}\}_{\ell=1}^{q}, \{\sigma_{\varepsilon,it}\}\right) &= \prod_{t=1}^{T} P\left(\varepsilon_{it} \middle| \{\theta_{i\ell}\}_{\ell=1}^{q}, \{\varepsilon_{it-\ell}\}_{\ell=1}^{t-1+q}, \{\sigma_{\varepsilon,it}\}\right) \\
&= \prod_{t=1}^{T} P\left(\varepsilon_{it} \middle| \{\theta_{i\ell}\}_{\ell=1}^{q}, \{v_{it-\ell}\}_{\ell=1}^{q}, \sigma_{\varepsilon,it}\right) \\
&= \prod_{t=1}^{T} \frac{1}{\sigma_{\varepsilon,it}} \phi\left(\frac{\varepsilon_{it} - \left(\sum_{\ell=1}^{q} \theta_{i\ell} v_{it-\ell}\right)}{\sigma_{\varepsilon,it}}\right)
\end{aligned}
$$

where ϕ is the standard normal density. This is the likelihood from a regression of ε_{it} on $(v_{it-1}, \ldots, v_{it-q})'$ with heteroskedastic Gaussian errors or, equivalently, (up to a scaling constant) from a regression of $y_{it} = \varepsilon_{it}/\sigma_{\varepsilon,it}$ on $x_{it} = (v_{it-1}, \ldots, v_{it-q})'/\sigma_{\varepsilon,it}$ with i.i.d. $N(0,1)$ errors. Since the prior is $(\theta_{i1}, \ldots, \theta_{iq})' \sim N(0_{Q\times 1}, V_\theta)$ where the variance is $V_\theta = \operatorname{diag}\left(v_1, \ldots, v_q\right)$, the posterior follows from the usual regression formula

$$
\{\theta_{i\ell}\}_{\ell=1}^{q} \middle| \{\varepsilon_{it}\}, \{\sigma_{\varepsilon,it}\} \sim N\left(\left[V_\theta^{-1} + \sum_{t=1}^{T} x_{it}x_{it}'\right]^{-1} \sum_{t=1}^{T} x_{it}y_{it}, \left[V_\theta^{-1} + \sum_{t=1}^{T} x_{it}x_{it}'\right]^{-1}\right).
$$

Conditioning on the initial observations $\{\varepsilon_{i1}, \ldots, \varepsilon_{i1-q}\}$ has at most a small effect when T is large. As an alternative, we can include $\{v_{it}\}$ as state variables in step

(A).

Implementation, numerical accuracy and tests We do $S = 12,000$ draws retaining one every two after burning the first $6,000$. The result is a chain for the parameters (θ, γ) with low enough autocorrelations that the posterior expectations have negligible Monte Carlo standard errors. We also monitor the behavior of the latent variables and stochastic volatilities, the paths of which seem to stabilize within a small region well before the burn-in period ends. We ran the posterior simulator test suggested by Geweke [2004] and an extensive Monte Carlo simulation study, finding no indication against our implementation.

C Monte Carlo simulations

To assess how reliable our estimates of Trend Wage Inflation are we conduct Monte Carlo simulations. We are interested in two questions. First, we ask whether our approach can accurately trace out the persistence pattern in monthly wage inflation from observations on 12-month wage growth rates, that is, whether we can successfully disentangle the temporal aggregation in the data. Second, we ask whether our approach has any bias — any tendency to over or understate the role of common and idiosyncratic components. To give a preview, the findings in this appendix validate the performance of the model in disentangling temporal aggregation and show that our method is not biased towards attributing an excessive role to the common component.

We simulate $n_{\mathrm{MC}} = 200$ samples of size $N = 7$ and $T = 300$ (N, T are chosen to be similar to our sample of wage growth by industry) from the following data generating process (DGP):

$$w_{it} = \frac{1}{12} \sum_{\ell=1}^{12} \tilde{\tau}_{i,t+1-\ell} + \tilde{\varepsilon}_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T,$$

$$
\begin{aligned}
\tilde{\tau}_{it} &= \alpha_{\tau,i}\tau_{ct} + \tau_{it}, \\
\tilde{\varepsilon}_{it} &= \alpha_{\varepsilon,i}\varepsilon_{ct} + \varepsilon_{it} + \theta_i \varepsilon_{i,t-1}, \\
\Delta\tau_{ct} &\overset{iid}{\sim} N(0,1), \quad \Delta\tau_{it} \overset{iid}{\sim} N(0,\sigma^2_{\Delta\tau,i}), \\
\varepsilon_{ct} &\overset{iid}{\sim} N(0,1), \quad \varepsilon_{it} \overset{iid}{\sim} N(0,\sigma^2_{\varepsilon,i}).
\end{aligned}
$$

We abstract from time-variation in loadings and volatilities in the DGP and we treat sectors symmetrically, setting $\alpha_{\tau,i} = \alpha_\tau$, $\alpha_{\varepsilon,i} = \alpha_\varepsilon$, $\sigma_{\Delta\tau,i} = \sigma_{\Delta\tau}$, $\sigma_{\varepsilon,i} = \sigma_\varepsilon$. Nonetheless, we conduct estimation in each sample allowing for both time-varying parameters and heterogeneity and we use exactly the same priors we adopted in the empirical analysis of the paper.

We calibrate $\alpha_\varepsilon = 0.02$, $\sigma_{\Delta\tau} = 0.2$ and $\sigma_\varepsilon = 0.85$ (we also set $\theta_i = \theta = 0.1$) using the averages across sectors and over time of the estimates we obtained in our sample of wage growth by industry. For α_τ we try two different values, namely: $\alpha_\tau \in \{0, 0.3\}$. Since the variance of $\Delta\tau_c$ is unity, the parameter α_τ controls the importance of the common component in driving the trajectory of the total trend of each sector. The value $\alpha_\tau = 0.3$ is the average across industries and periods we found in our sample. The value $\alpha_\tau = 0$ represents a case where the common component is zero. We consider this extreme case to assess whether our method would spuriously recover a common component that does not exist.

In each sample, we run our estimation algorithm and we recover the posterior p-quantile of the total trend $\tilde{\tau}_t = N^{-1}\sum_{i=1}^N \tilde{\tau}_{it}$, its common part $\tilde{\tau}_{Ct} = N^{-1}\sum_{i=1}^N \alpha_{\tau,i}\tau_{ct}$ and its idiosyncratic part $\tilde{\tau}_{It} = N^{-1}\sum_{i=1}^N \tau_{it}$, that we denote $\tilde{\tau}_t(p)$, $\tilde{\tau}_{Ct}(p)$ and $\tilde{\tau}_{It}(p)$. Note that we are assuming all sectors have the same employment share and that these are constant over time, i.e., we set $s_{it} = N^{-1}$.

We use our Monte Carlo simulation to estimate the bias of the posterior median (seen as a point estimate of the latent variables) and the frequentist coverage rates of credible intervals based on the posterior. For $\tilde{\tau}_t$, for example, we have

$$\text{bias}_t = E\left[\tilde{\tau}_t\left(\frac{1}{2}\right) - \tilde{\tau}_t\right],$$
$$\text{cov}_t = P\left[\tilde{\tau}_t\left(\frac{\beta}{2}\right) \leq \tilde{\tau}_t \leq \tilde{\tau}_t\left(1 - \frac{\beta}{2}\right)\right]$$

where expectations and probabilities are taken with respect to repeated sampling from the DGP (and they are estimated by averaging over the n_{MC} Monte Carlo samples). A good estimator of the trend will deliver $\text{bias}_t \approx 0$ and $\text{cov}_t \approx 1 - \beta$. We can similarly define bias and coverage rates for $\tilde{\tau}_{Ct}$ and $\tilde{\tau}_{It}$.

One detail is that the location of $\tilde{\tau}_{Ct}$ and $\tilde{\tau}_{It}$ has to be decided by a normalization. In the empirical analysis of the paper, for example, we use $\tilde{\tau}_{C1} = 0$. To avoid ambiguities, below we report bias and coverage rates for $\tilde{\tau}_{Ct} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Cs}$ and $\tilde{\tau}_{It} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Is}$. Results look similar using alternative normalizations.

We display the bias calculations in Figure C1. For $\tilde{\tau}_t$, for example, we plot the sampling distribution of $\left\{\tilde{\tau}_t\left(\frac{1}{2}\right) - \tilde{\tau}_t\right\}_{t=1}^{T}$ indicating for each t the values contained between the 0.16- and 0.84-quantiles of the sampling distributions with a shaded area. We also report $\text{med}\left(\tilde{\tau}_t\left(\frac{1}{2}\right) - \tilde{\tau}_t\right)$ (blue dashed line) and $\text{bias}_t = E\left[\tilde{\tau}_t\left(\frac{1}{2}\right) - \tilde{\tau}_t\right]$ (black dotted line). We do the same for $\tilde{\tau}_{Ct}$ and $\tilde{\tau}_{It}$. The figure shows that our approach has no systematic tendency to over or underestimate the trend, its common or its idiosyncratic component. This holds for both $\alpha_\tau = 0.3$ (a value representative of our sample) and, reassuringly, for $\alpha_\tau = 0$. In other words, even in the extreme case where the common component does not exist, there is no evidence to suggest that our model would spuriously find a role for a common component.

Turning to the coverage properties of posterior intervals, the performance of our method is solid. We report the average over t of estimated coverage rates $T^{-1}\sum_{t=1}^{T}\text{cov}_t$ for our two designs in Table C1.

We set the probability level to $1 - \beta = 0.68$, the level we use in our empirical analysis

FIGURE C1. Bias of the posterior median estimate

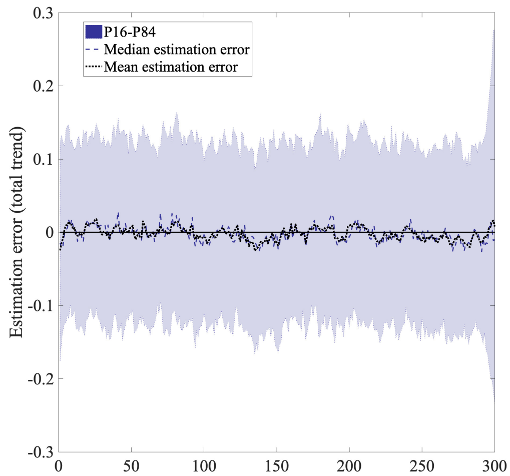

(a) Bias for $\tilde{\tau}_t$ for $\alpha_\tau = 0$

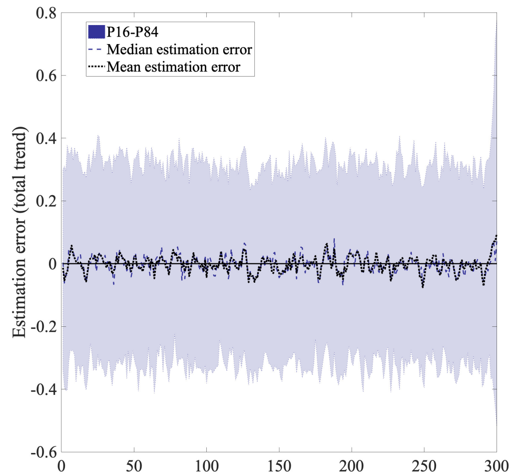

(b) Bias for $\tilde{\tau}_t$ for $\alpha_\tau = 0.3$

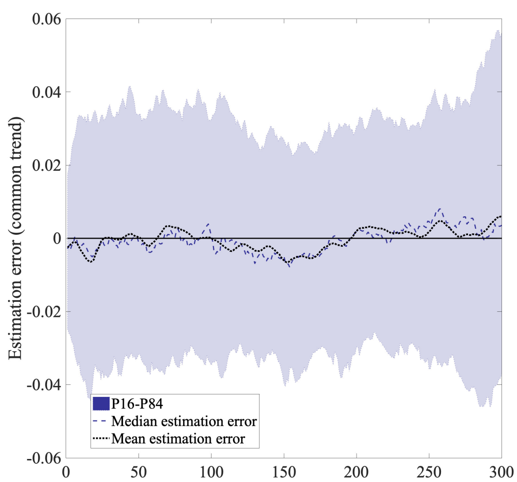

(c) Bias for $\tilde{\tau}_{Ct} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Cs}$ for $\alpha_\tau = 0$

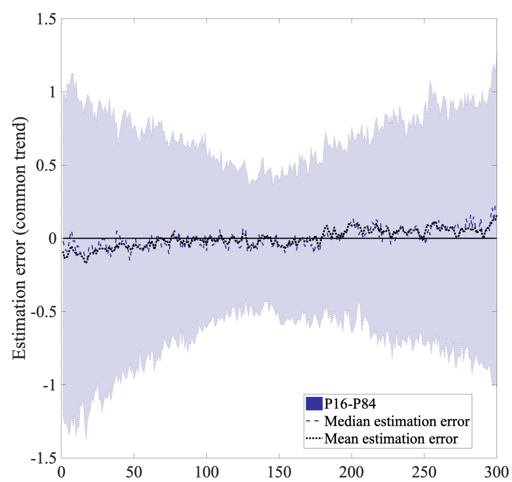

(d) Bias for $\tilde{\tau}_{Ct} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Cs}$ for $\alpha_\tau = 0.3$

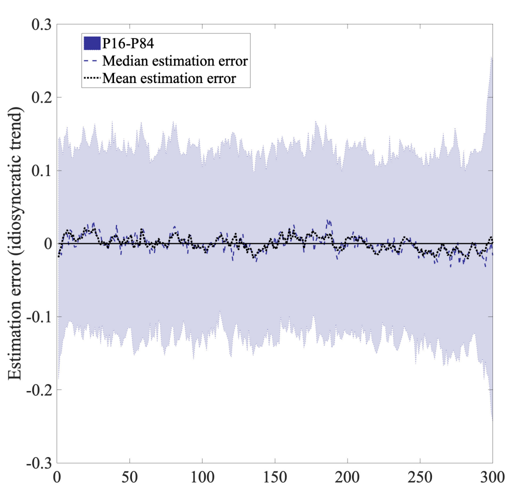

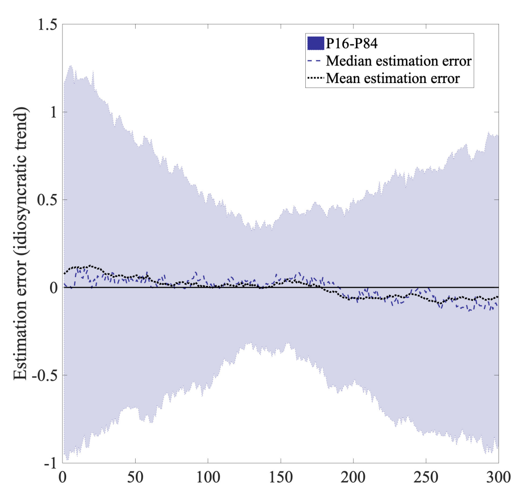

(e) Bias for $\tilde{\tau}_{It} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Is}$ for $\alpha_\tau = 0$

(f) Bias for $\tilde{\tau}_{It} - T^{-1}\sum_{s=1}^{T}\tilde{\tau}_{Is}$ for $\alpha_\tau = 0.3$

TABLE C1. Average coverage rates for nominal rate $1-\beta = 0.68$

	$\alpha_\tau = 0$	$\alpha_\tau = 0.3$
$\tilde{\tau}_t$	0.763	0.719
$\tilde{\tau}_{Ct}$	0.998	0.665
$\tilde{\tau}_{It}$	0.896	0.657

and equivalent to intervals of roughly one standard deviation radius under a normal distribution. When $\alpha_\tau = 0.3$, the average coverage rates are reasonably close to the nominal rate suggesting that our framework produces reliable inferences about the trend and its common and idiosyncratic components in repeated samples.[3]

When $\alpha_\tau = 0$, our method produces relatively conservative inferences in the sense that it overcovers both the common and idiosyncratic component. In particular, the probability bands for $\tilde{\tau}_{Ct}$ contain the zero line (its true value) in practically all samples. This agrees with our claim that our method does not have a systematic tendency to find a common component when there is none.

These results are important because the good coverage of our method is a frequentist property, even though the intervals we use are Bayesian credible intervals. Moreover, our estimation and inference approach uses a prior that does not center the model at the DGP, suggesting that shrinking our model away from the DGP has a negligible effect as with these parameters and sample size the prior is dominated by the sample information.

[3]Our method also achieves good coverage pointwise in t.

D Additional empirical results

D.1 Wage growth across industries

The figures in this section show the raw data, common, and total trend for the seven industries in our baseline specification. We also report outlier probabilities for each time period. These probabilities are generally low, typically below one third. However, they are informative for specific episodes. For instance, the model attaches a 15% probability that the increase in wage growth in Leisure and Hospitality, in August 2021, was an outlier. This stands in sharp contrast with the high readings between December 2021 and June 2022, which are not measured as likely outliers.

FIGURE D1. Total and common trend by industry

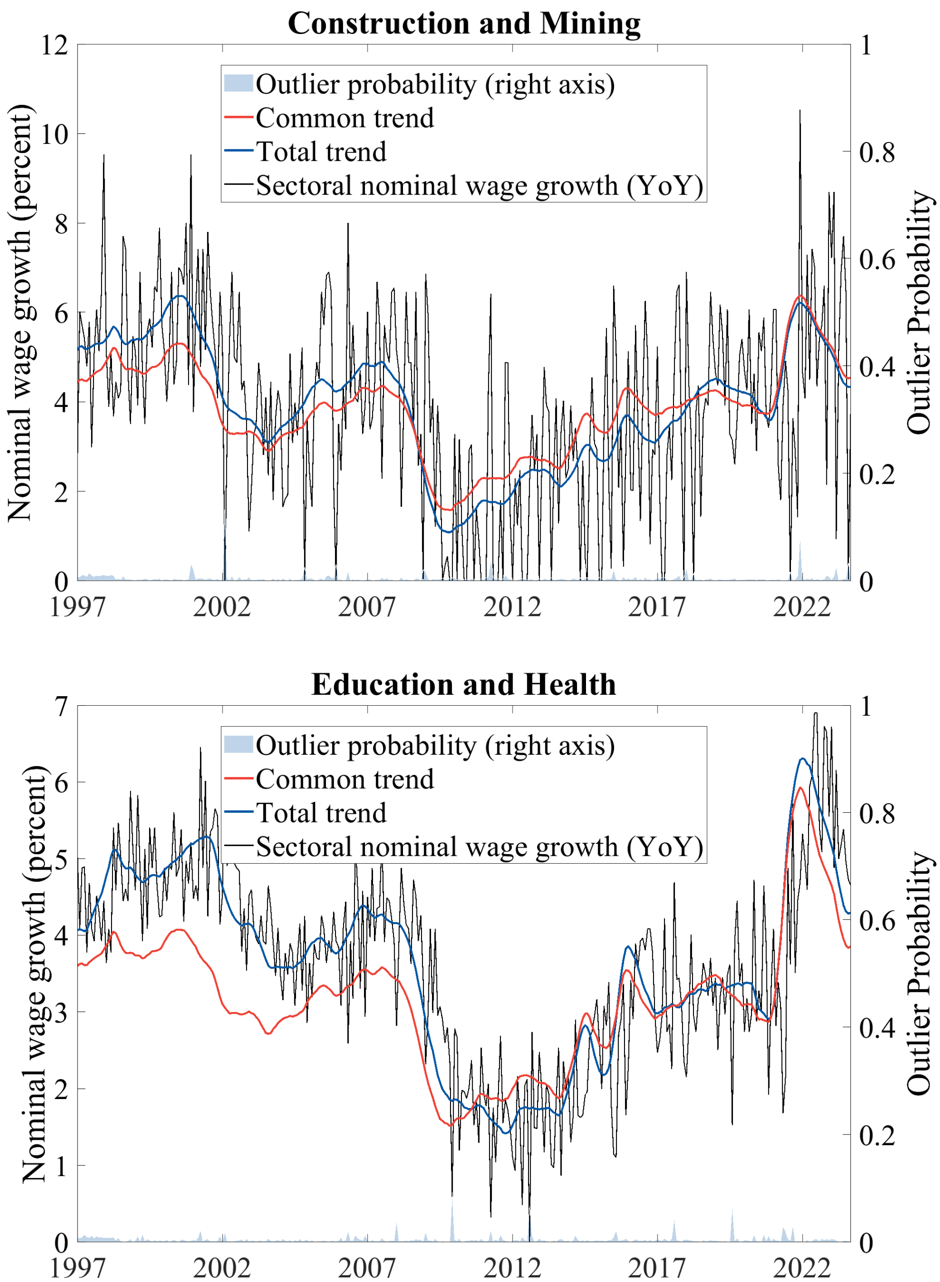

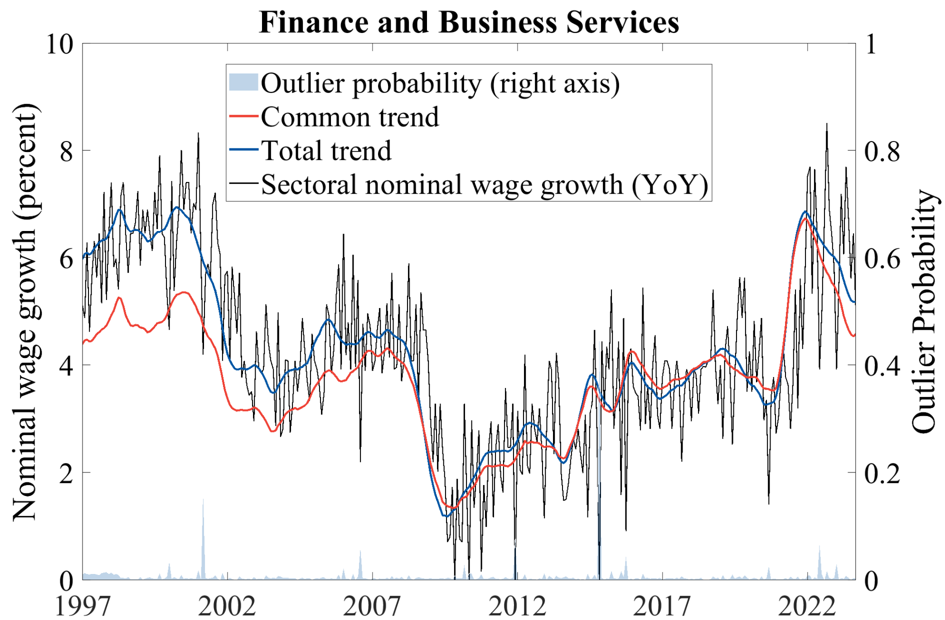
Finance and Business Services
Outlier probability (right axis)
Common trend
Total trend
Sectoral nominal wage growth (YoY)
Nominal wage growth (percent)
Outlier Probability
10
8
6
4
2
0
1
0.8
0.6
0.4
0.2
0
1997
2002
2007
2012
2017
2022

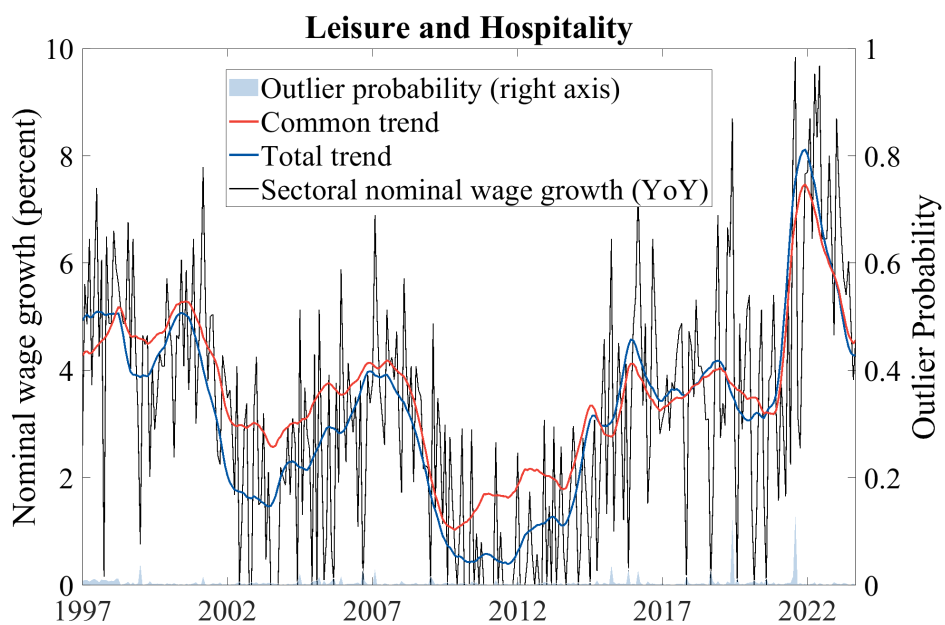
Leisure and Hospitality
Outlier probability (right axis)
Common trend
Total trend
Sectoral nominal wage growth (YoY)
Nominal wage growth (percent)
Outlier Probability
10
8
6
4
2
0
1
0.8
0.6
0.4
0.2
0
1997
2002
2007
2012
2017
2022

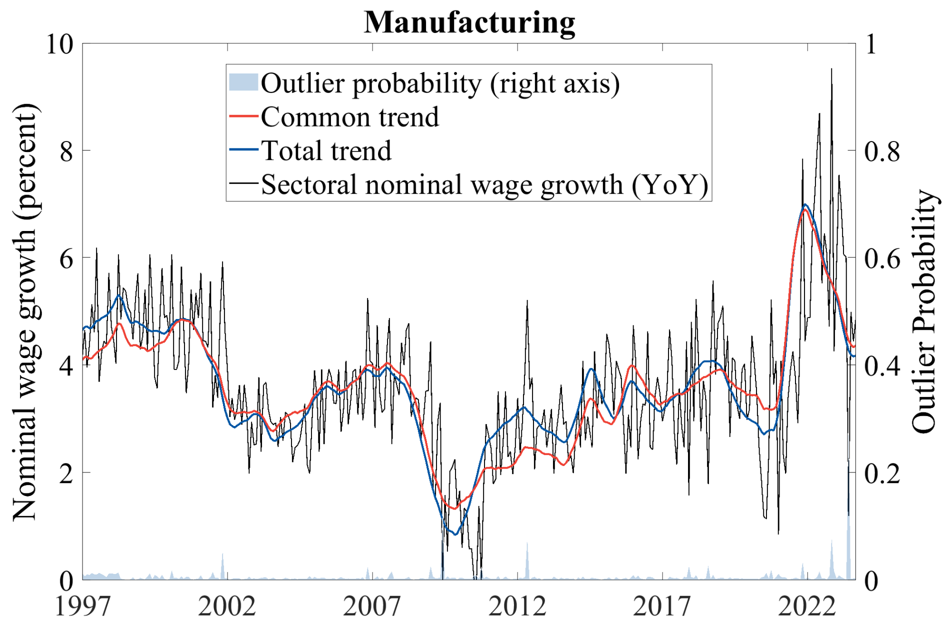
Manufacturing
Outlier probability (right axis)
Common trend
Total trend
Sectoral nominal wage growth (YoY)
Nominal wage growth (percent)
Outlier Probability
0
2
4
6
8
10
0
0.2
0.4
0.6
0.8
1
1997
2002
2007
2012
2017
2022

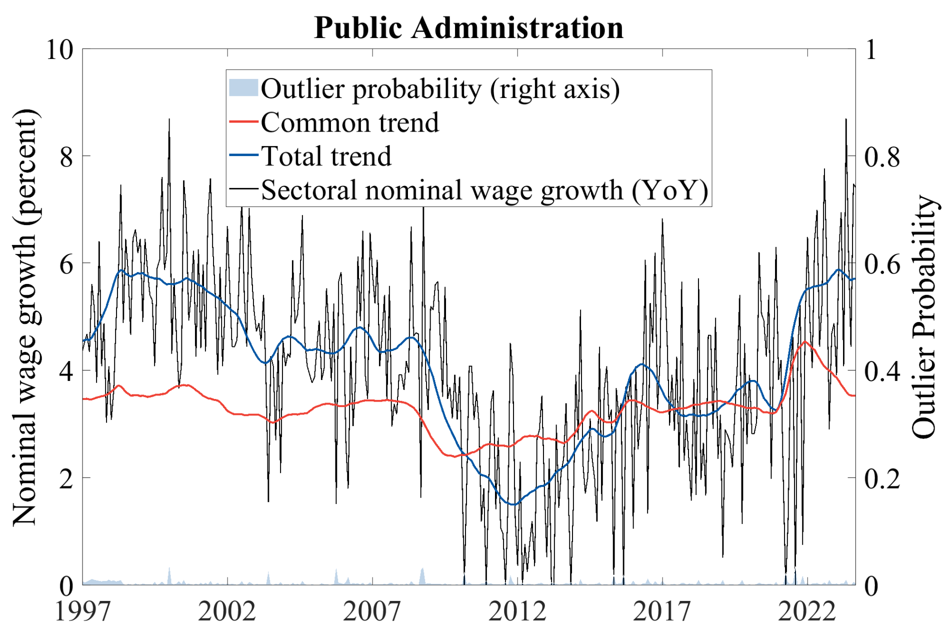
Public Administration
Outlier probability (right axis)
Common trend
Total trend
Sectoral nominal wage growth (YoY)
Nominal wage growth (percent)
Outlier Probability
0
2
4
6
8
10
0
0.2
0.4
0.6
0.8
1
1997
2002
2007
2012
2017
2022

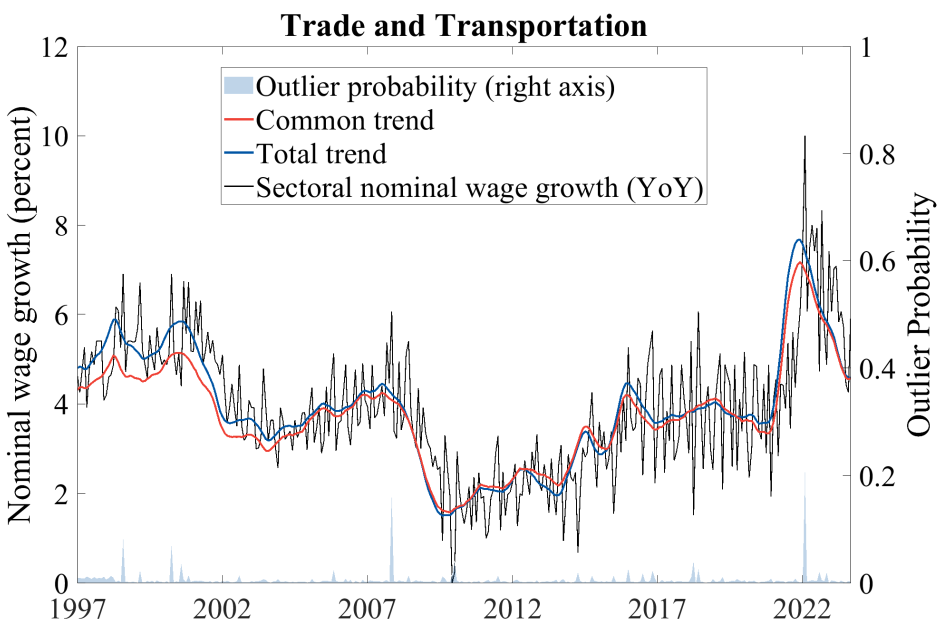

NOTES. The figure shows for each industry the raw nominal wage growth data, the total trend ($\alpha_{\tau,it}\tau_{ct} + \tau_{it}$), the common trend component ($\alpha_{\tau,it}\tau_{ct}$), and the outlier probability over the sample period.

D.2 Time-varying parameters

FIGURE D2. Estimates of time-varying parameters

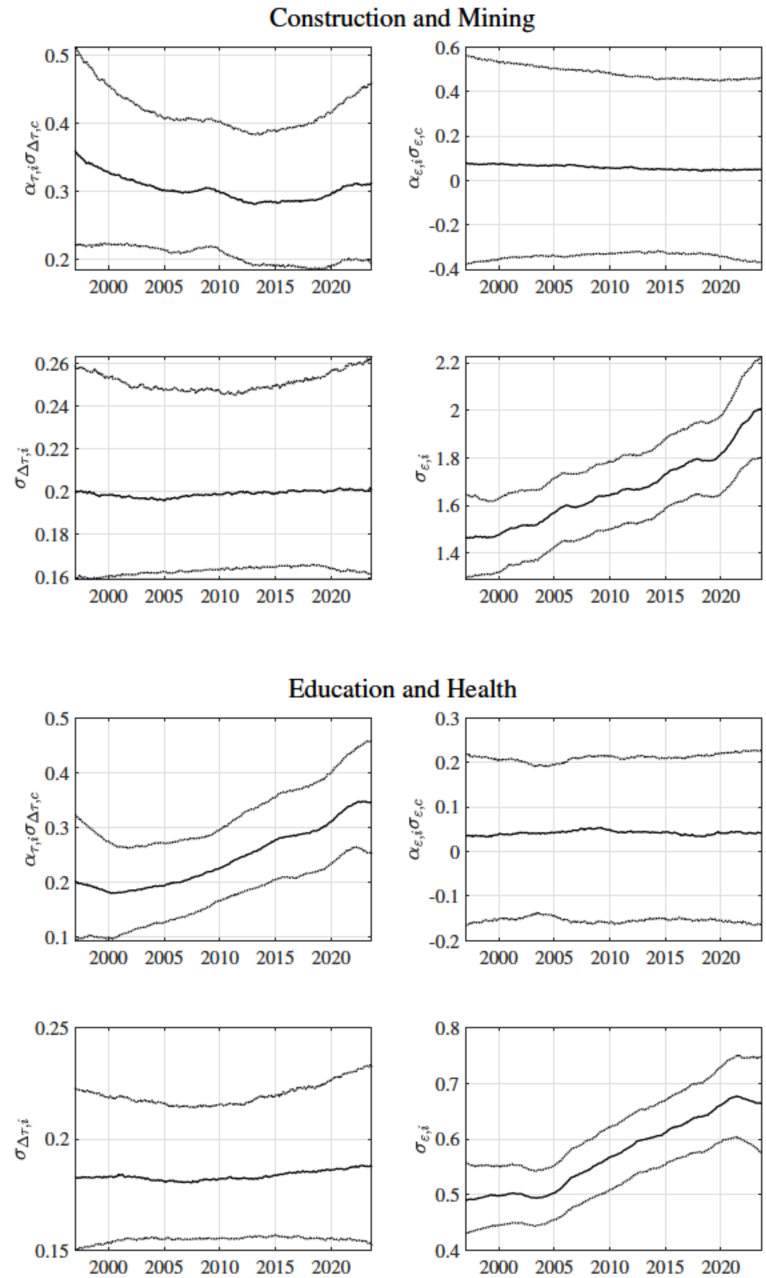

NOTES. Each panel reports the 0.16-quantile (lower/dotted line), 0.50-quantile (middle/solid line) and 0.84-quantile (upper/dotted line) from the posterior distribution of time-varying parameters at each time t. We report the product of loadings and standard deviations of common components because their scales are not separately pinned down.

FIGURE D3. Estimates of time-varying parameters

Finance and Business Services

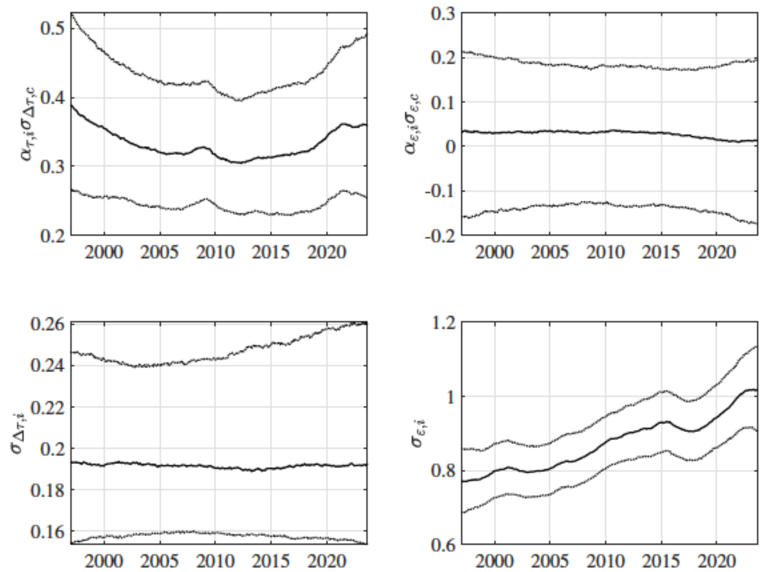

Leisure and Hospitality

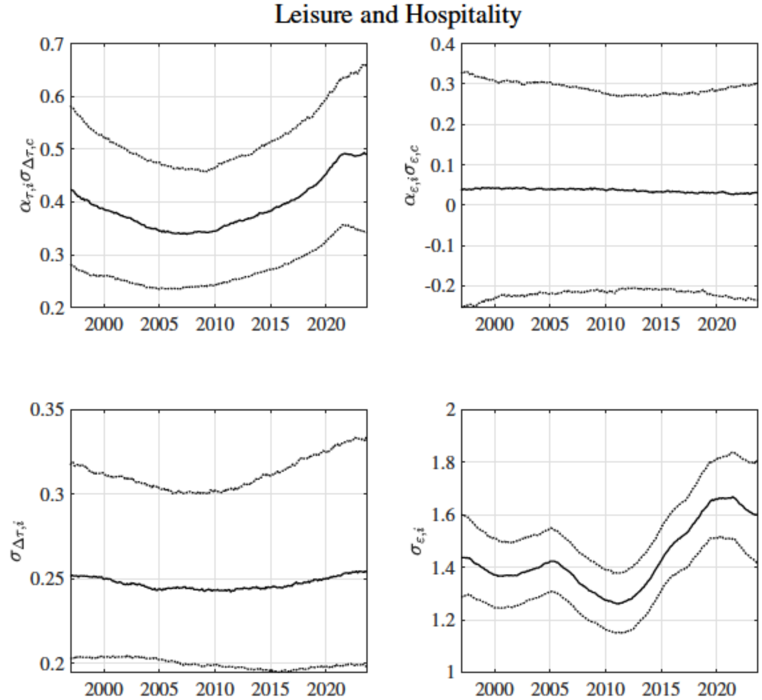

NOTES. Each panel reports the 0.16-quantile (lower/dotted line), 0.50-quantile (middle/solid line) and 0.84-quantile (upper/dotted line) from the posterior distribution of time-varying parameters at each time t. We report the product of loadings and standard deviations of common components because their scales are not separately pinned down.

FIGURE D4. Estimates of time-varying parameters

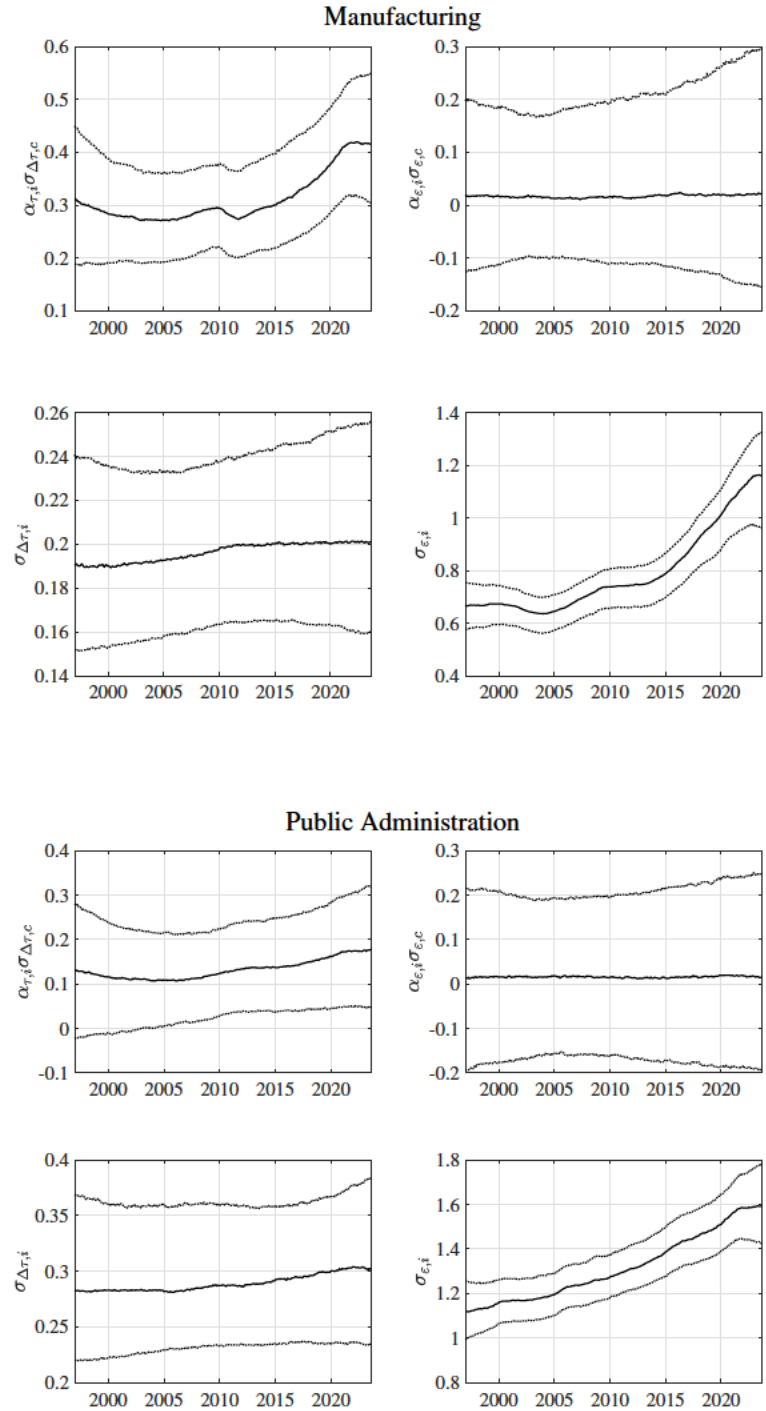

NOTES. Each panel reports the 0.16-quantile (lower/dotted line), 0.50-quantile (middle/solid line) and 0.84-quantile (upper/dotted line) from the posterior distribution of time-varying parameters at each time t. We report the product of loadings and standard deviations of common components because their scales are not separately pinned down.

FIGURE D5. Estimates of time-varying parameters

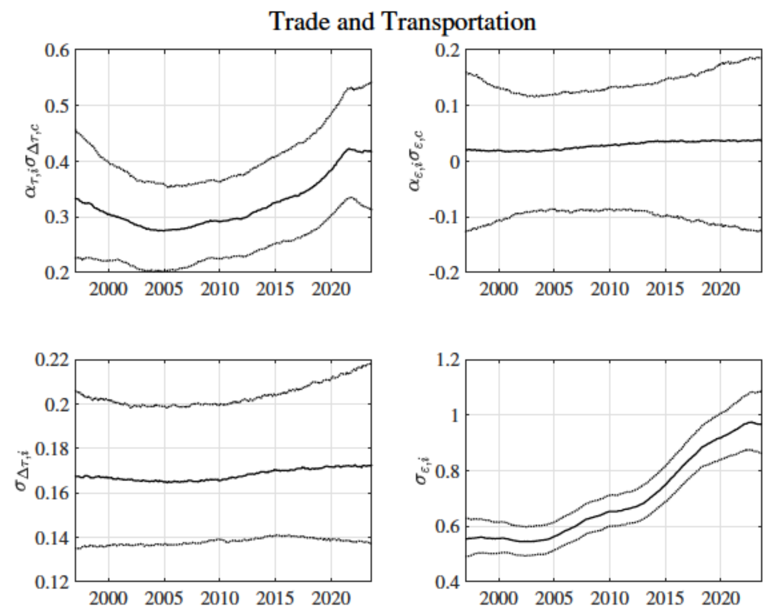

NOTES. Each panel reports the 0.16-quantile (lower/dotted line), 0.50-quantile (middle/solid line) and 0.84-quantile (upper/dotted line) from the posterior distribution of time-varying parameters at each time t. We report the product of loadings and standard deviations of common components because their scales are not separately pinned down.

D.3 Trend Wage Inflation using alternative CPS partitions

FIGURE D6. Trend Wage Inflation across models

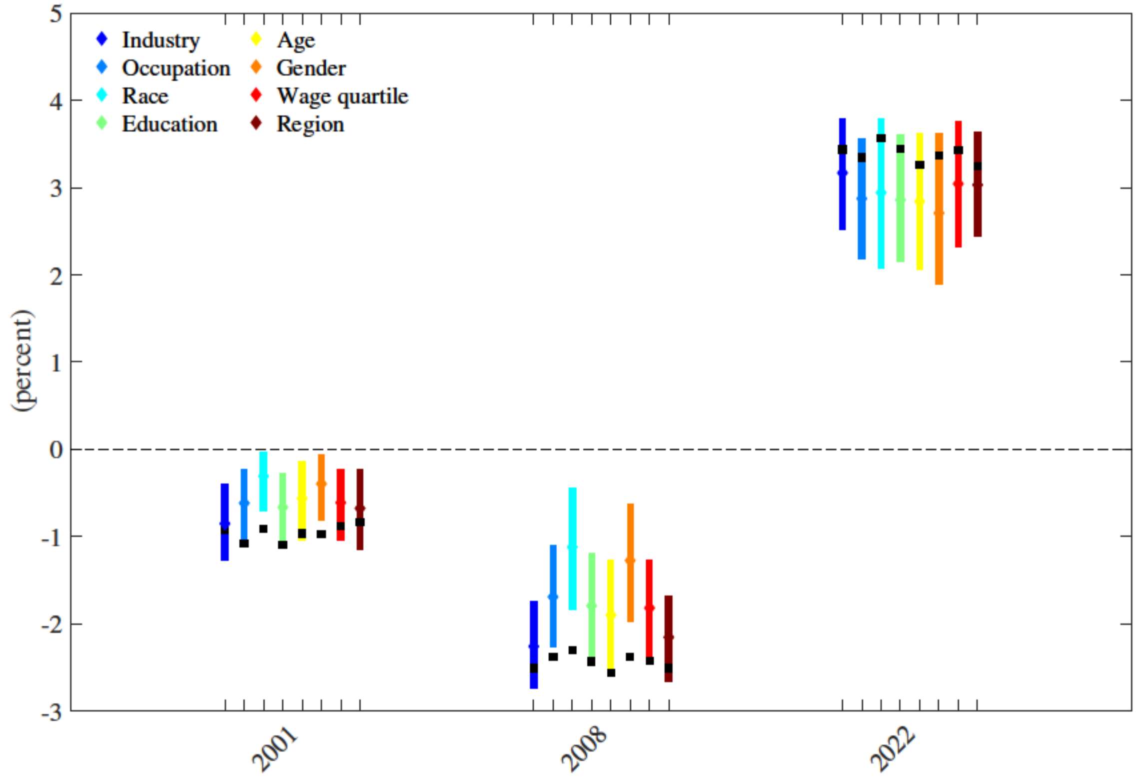

NOTES. The three episodes refer to the following periods: 2001m3-2001m11, 2007m12-2009m6, and 2020m6-2022m2. Each model uses the data cut described in the legend; the variables follow the Atlanta Fed Wage Growth tracker definition and are detailed in Appendix A. Black square markers indicate the peak-to-trough change in Trend Wage Inflation for each model and episode. The diamond markers indicate the peak-to-trough change for the common component $\left(\sum_{i=1}^{n} s_{it}\alpha_{\tau,it}\right)\tau_{ct}$ in each model and episode. Vertical lines show the 68 percent probability bands.

E Robustness checks

In this appendix, we verify the robustness of the main results of the paper to three choices we make in the empirical analysis. The first is to use the median instead of the mean of year-over-year wage growth as the observable w_{it} in our model. The second choice is to use the unweighted median as opposed to the median weighted by the survey weights as w_{it}. Third, we do not allow for $\tilde{\tau}_{it}$ to be itself the sum of a persistent and a transitory component. Figures E1, E2 and E3 show that both the historical behavior of the persistent component in wage growth and the relatively high importance of common variation across industries are insensitive to these choices.

FIGURE E1. Estimates based on mean year-over-year wage growth

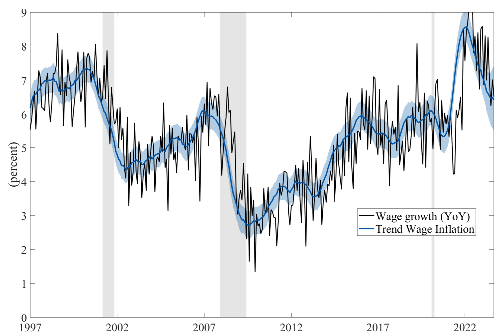

(a) Persistent component of wage growth

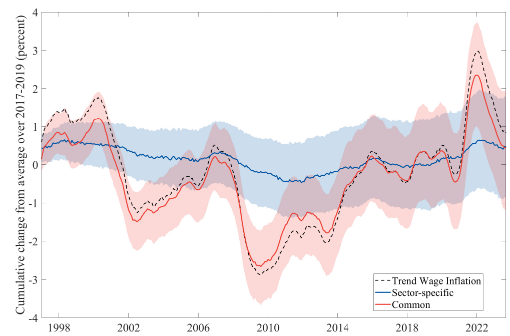

(b) Common and sector-specific contribution

Despite mean year-over-year wage growth being more volatile than median wage growth, our model traces a remarkably similar historical evolution of the persistent component (Trend Wage Inflation), with the largest swings located around the same episodes we discussed in Section 4 (i.e., the 2001 and 2008 recessions, and the post-pandemic inflation spike). Trend Wage Inflation is somewhat higher when using the mean instead of the median due to the positive skewness in the wage growth distribution, but this seems to imply merely a level shift in the persistent component. The cumulative changes in panel (b) of figure E1, for example, are quantitatively very close to our baseline results.

Differences in our estimates when using the unweighted instead of the weighted median of wage growth as w_{it} are imperceptible, as shown in figure E2.

FIGURE E2. Estimates based on unweighted median wage growth

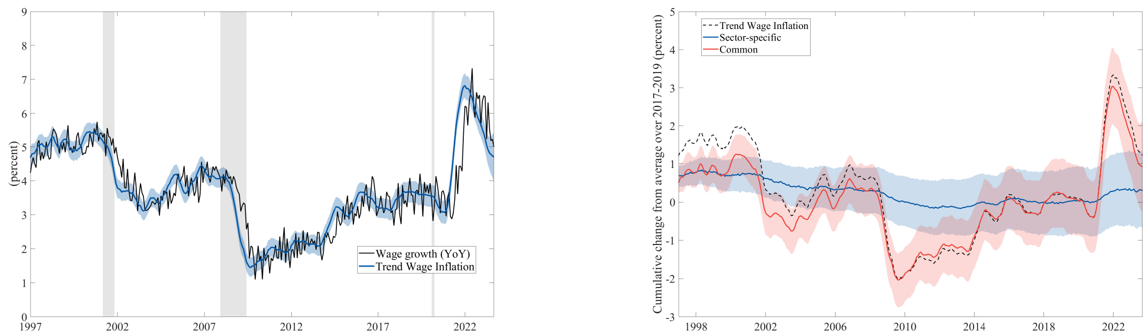

(a) Persistent component of wage growth

(b) Common and sector-specific contribution

Figure E3 illustrates a point made in section 3. Our empirical analysis interprets the transitory component of year-over-year wage growth $\tilde{\varepsilon}_{it}$ as being largely measurement error. Therefore, $\tilde{\tau}_{it}$ is interpreted as the unobservable monthly growth rate of nominal wages that could be recovered with a perfect error-free survey. Because we rely on time series smoothing techniques, the assumption that $\tilde{\tau}_{it}$ is well approximated by a random walk is important in order to filter the survey measurement error out. If instead $\tilde{\tau}_{it}$ is the sum of two components,

$$\tilde{\tau}_{it} = \tilde{\tau}_{it}^{\text{pers}} + \tilde{\tau}_{it}^{\text{tr}}$$

where $\tilde{\tau}_{it}^{\text{pers}}$ is now a random walk and $\tilde{\tau}_{it}^{\text{tr}}$ is white noise, our baseline model with the choice of moving average orders $p = q = 12$ would estimate $\tilde{\tau}_{it}^{\text{pers}}$ instead of $\tilde{\tau}_{it}$. Comparing estimates from this extended model and the results in section 4 provides a sense of how important the genuine transitory shock $\tilde{\tau}_{it}^{\text{tr}}$ is. Figure E3 indicates that $\tilde{\tau}_{it}^{\text{tr}}$ plays at most a minor role and that the more parsimonious model used in our paper captures sufficiently well the most salient movements in aggregate wage growth, which tend to be very persistent.

A final piece of evidence supporting our interpretation of $\tilde{\varepsilon}_{it}$ as measurement error is

FIGURE E3. Estimates based on a more flexible specification

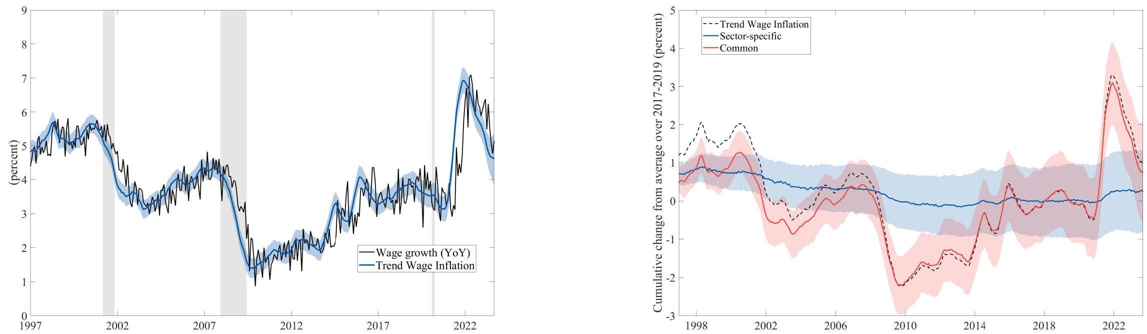

(a) Persistent component of wage growth

(b) Common and sector-specific contribution

shown in figure E4. Consider the transitory component of aggregate wage growth, which we define as

$$\tilde{\varepsilon}_t = \sum_{i=1}^{n} s_{it}\tilde{\varepsilon}_{it} = \sum_{i=1}^{n} s_{it}w_{it} - \tilde{\tau}_t$$

where s_{it} is the employment share of cross-section i in month t. The variance of $\tilde{\varepsilon}_t$ is given by

$$\tilde{\sigma}^2_{\varepsilon,t} = \left(\sum_{i=1}^{n} s_{it}\alpha_{\varepsilon,it}\right)^2 \sigma^2_{\varepsilon,ct} + \sum_{i=1}^{n} s^2_{it}\sigma^2_{\varepsilon,it}.$$

If $\tilde{\varepsilon}_{it}$ is the error made in using the sample median of year-over-year wage growth w_{it} from a sample of n_{it} workers to estimate the population growth rate $\sum_{\ell=1}^{12} \tilde{\tau}_{it+1-\ell}/12$ in each sector i, then the standard deviation $\tilde{\sigma}_{\varepsilon,t}$ should be proportional to $1/\sqrt{n_t}$ where $n_t = \sum_{i=1}^{n} n_{it}$ is the survey sample size in month t. Figure E4 shows precisely that: a scatter plot of (the posterior median estimate of) $\tilde{\sigma}_{\varepsilon,t}$ against n_t in which most of the points lie close to the line $\tilde{\sigma}_{\varepsilon,t} = \hat{c}/\sqrt{n_t}$.[4] We find a similar pattern if we consider the correlation between sample size n_{it} and the standard deviation $\sigma_{\varepsilon,it}$ for a specific industry i.

We also find, consistent with our interpretation, that for every i the path of $\alpha_{\varepsilon,it}\sigma^2_{\varepsilon,ct}$

[4] In fact, the correlation between $\tilde{\sigma}_{\varepsilon,t}$ and $1/\sqrt{n_t}$ is 0.9.

FIGURE E4. Standard deviation of transitory shocks and survey sample sizes

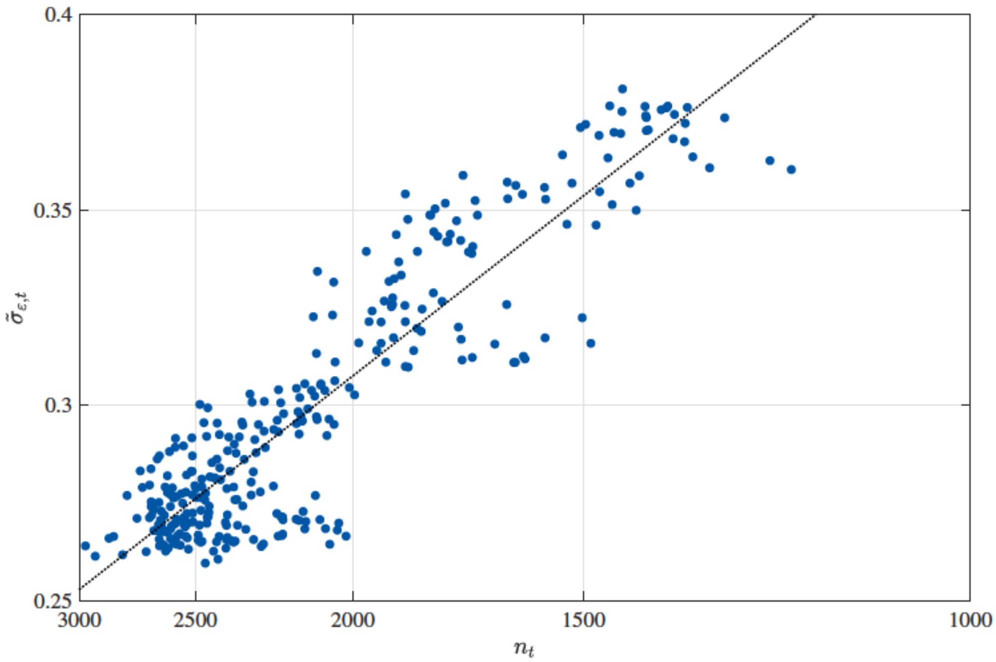

always contains the zero line, indicating a negligible role for cross-sectional correlation across $\tilde{\varepsilon}_{it}$.

F Additional evidence using CES

This appendix presents estimates of the persistent component of month-on-month growth rates in nominal wages using data from the Current Employment Statistics (CES).[5] Because the data already provides month-on-month changes, denoted by W_{it} below, we estimate our model without temporal aggregation. In other words, instead of (2), our measurement equation is

$$W_{it} = \tilde{\tau}_{it} + \tilde{\varepsilon}_{it}$$

[5]We use average hourly earnings of production and non-supervisory employees on private nonfarm payrolls.

with the persistent component $\tilde{\tau}_{it}$ and the transitory component $\tilde{\varepsilon}_{it}$ modeled as in Section 3. The cross-sectional dimension is industries since the CES is a survey of establishments.[6]

As noted in Section 2, the CES measure of wages is subject to compositional issues. However, the CES spans a longer period (in this case beginning in 1964), which allows us to empirically study additional recessions and the inflationary episodes of the late 1960s and 1970s. Figure F1 plots the trend estimates against the raw data, while the decomposition into common and sector-specific drivers is shown in Section 4.3.

Figure F1 shows that the model attributes most of the high-frequency variation in nominal wage growth in the CES to the transitory variation term $\tilde{\varepsilon}_{it}$. The two largest changes in the persistent component of wage inflation correspond to the inflation episodes in the 1970s and the post-pandemic inflation surge. From the 1980s, most NBER recessions tend to be associated with a drop in Trend Wage Inflation.

FIGURE F1. Estimates using CES data

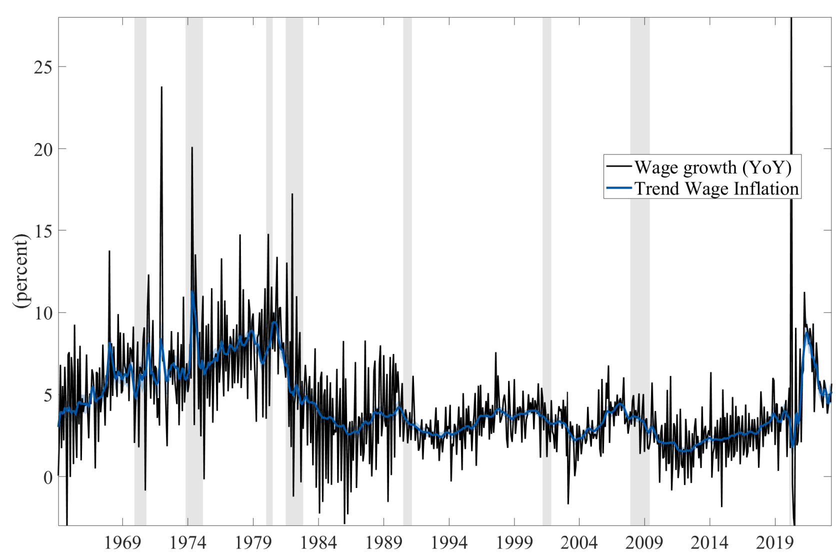

[6] We consider 10 industries: Construction, Financial Activities, Information, Leisure and Hospitality, Manufacturing, Mining and Logging, Other Services, Private Education and Health Services, Professional and Business Services, Trade-Transportation-Utilities.

G Estimates using a longer CPS sample with missing data

The three panels of Figure G1 compare the posterior median of trend components based on a longer sample with missing data (from 1983m1 to 2023m9; blue solid line) and a shorter sample with no missing data (1997m1 to 2023m9; red dotted line).

For the longer sample there are two gaps of missing data of 18 months each (one beginning around 1985 and another around 1995). In order to compute the aggregate trends over those periods we linearly interpolate the sectoral shares at the end points. We argue this is a reasonable approach as there is nothing to suggest there was a large (and fast) sectoral relocation during those intervals.

Reassuringly, the estimates on the overlapping period are very close to each other and display the same dynamics over time. Moreover, the longer sample allows us to incorporate an additional recession in 1990 where we clearly see the large drop in the persistent component of wage growth driven by the common component across sectors, reminiscent of the recessions we had in our shorter sample.

FIGURE G1. Estimates based on the 1983m1-2023m9 and 1997m1-2023m9 samples

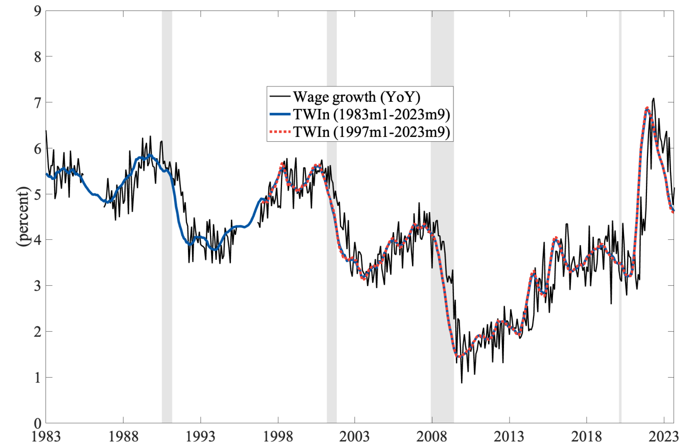

(a) Trend wage inflation

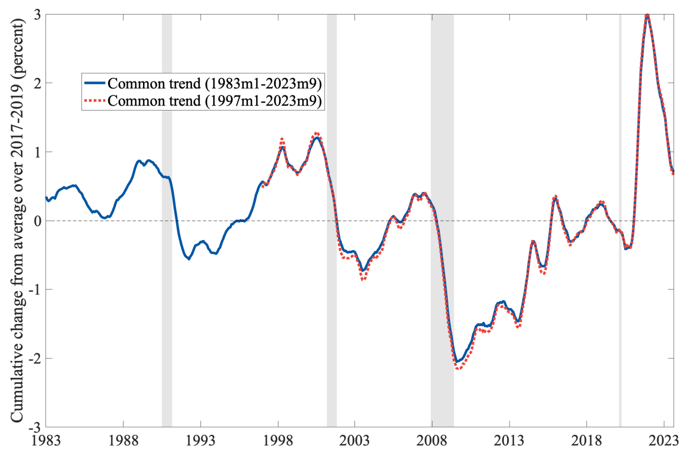

(b) Common component of trend wage inflation

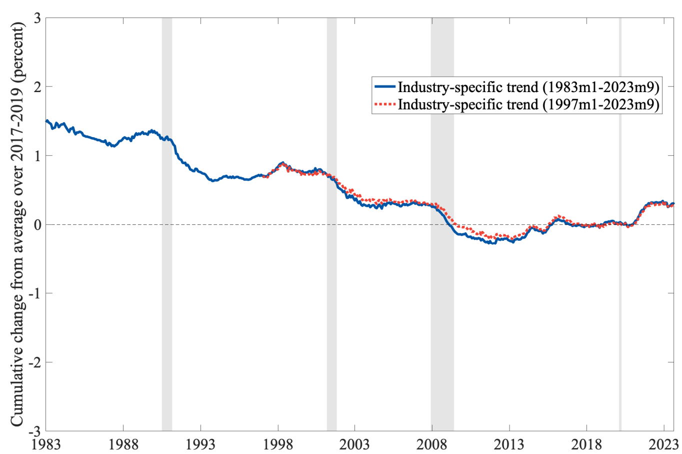

(c) Sector-specific component of trend wage inflation

References: Supplemental Appendix

M. Del Negro and C. Otrok. Dynamic factor models with time-varying parameters: measuring changes in international business cycles. Technical report, 2008.

T. Doan, R. B. Litterman, and C. A. Sims. Forecasting and Conditional Projection Using Realistic Prior Distributions. NBER Working Papers 1202, National Bureau of Economic Research, Inc, Sept. 1983. URL https://ideas.repec.org/p/nbr/nberwo/1202.html.

J. Durbin and S. J. Koopman. A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89(3):603–615, 2002.

J. Geweke. Getting it right: Joint distribution tests of conditional simulators. *Journal of the American Statistical Association*, 99:799–804, 2004.

A. Inoue and L. Kilian. Joint confidence sets for structural impulse responses. *Journal of Econometrics*, 192(2):421–432, 2016. doi: 10.1016/j.jeconom.2016.02. URL https://ideas.repec.org/a/eee/econom/v192y2016i2p421-432.html.

S. Kim, N. Shephard, and S. Chib. Stochastic volatility: Likelihood inference and comparison with arch models. *Review of Economic Studies*, 65:361–393, 1998.

Y. Omori, S. Chib, N. Shephard, and J. Nakajima. Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140:425–449, 2007.

J. H. Stock and M. W. Watson. Core inflation and trend inflation. *Review of Economics and Statistics*, 98(4):770–784, 2016.