

## Downward Rigidity in the Wage for New Hires<sup>†</sup>

By JONATHON HAZELL AND BLEDI TASKA\*

*Wage rigidity is an important explanation for unemployment fluctuations. In benchmark models wages for new hires are key, but there is limited evidence on this margin. We use wages posted on vacancies, with job and establishment information, to measure the wage for new hires. We show that our measure of the wage for new hires is rigid downward and flexible upward, in two steps. First, wages change infrequently at the job level, and fall especially rarely. Second, wages do not respond to rises in unemployment, but respond strongly to falls in unemployment. Job information is crucial for detecting downward rigidity. (JEL E24, E32, J23, J31, J63, M51)*

Suppose there is downward wage rigidity, meaning that wages do not fall during recessions. Economists have long argued that unemployment should then rise, because the cost of labor remains high even as labor demand falls (Keynes 1936). Downward wage rigidity for *new hires* is particularly important, because employment is a long term contract. Therefore the present value of wages matters to workers and firms (Barro 1977). One important component of the present value of wages is the wage for new hires (Pissarides 2009). Even if wages in continuing jobs change little, the present value of wages can still vary if the wage for new hires is flexible.

There is not yet consensus about wage rigidity for new hires. A key difficulty is job composition. Pissarides (2009) surveys some work on the wage for new hires using worker-level survey data without job information. This work studies the average wage for new hires, controlling for worker characteristics but averaging across the jobs that hire these workers. However, if job composition varies over time, then average wage changes reflect changing job composition, as well as wage changes

\* Hazell: London School of Economics (email: [j.hazell@lse.ac.uk](mailto:j.hazell@lse.ac.uk)); Taska: Burning Glass Technologies (email: [bt540@nyu.edu](mailto:bt540@nyu.edu)). Arnaud Costinot was the coeditor for this article. Hazell thanks Iván Werning, Daron Acemoglu, Jonathan Parker and Emi Nakamura for invaluable support and guidance; and Seyed Mahdi Hosseini, Marina Feliciano and Borui Zhu for superb research assistance. We also thank for their comments the coeditor, four anonymous referees, Martin Beraja, Sydnee Caldwell, Gabriel Chodorow-Reich, Allie Cole, Tyler Cowen, Maarten De Ridder, Juliette Fournier, John Grigsby, Erik Hurst, Ethan Ilzetzki, Gwen Jing, Mazi Kazemi, Eben Lazarus, Chen Lian, Benjamin Moll, Christian Moser, Christina Patterson, Ricardo Reis, Ishaana Talesara, Claudia Sahm, Pari Sastry, Anna Stansbury, Jón Steinsson, Emil Verner, Arthur Wickard and Sammy Young; as well as participants at the 2018 NBER Summer Institute, Monetary Economics, the MIT Macro and Labor lunches; and seminars at the London School of Economics, Chicago Booth School of Business, Columbia Business School, the University of Michigan, Brown University, UCLA Anderson School of Management, the University of Cambridge, CREI, Yale and Northwestern. Hazell thankfully acknowledges funding from the Becker-Friedman Institute at the University of Chicago; the Washington Center for Equitable Growth; the Kennedy Memorial Trust; and the Julis-Rabinowitz Center for Public Policy and Finance and the Griswold Center for Economic Policy at Princeton University. This paper previously circulated under the title “Posted Wage Rigidity.”

<sup>†</sup>Go to <https://doi.org/10.1257/aer.20201793> to visit the article page for additional materials and author disclosure statement(s).

for individual jobs (Gertler, Huckfeldt, and Trigari 2020). As an example, consider an economy of high-wage bankers and low-wage baristas, and suppose the share of barista hires increases. Then, average wages for new hires may fall, even if wages fall for neither baristas nor bankers. Equally, if the share of barista hires decreases, then average wages might not fall and could even rise, even if wages fall for both baristas and bankers.

To confront the difficulty of job composition, this paper studies a measure of the wage for new hires with job information. Our dataset, provided by Burning Glass Technologies, contains wages posted in online vacancies, with job titles, establishment identifiers, and pay frequency, for the United States. The source is vacancies from online job boards and company websites. Our main sample covers January 2010 to June 2020, starting after the Great Recession and finishing after the contractionary phase of the COVID-19 recession. We show that posted wages are a reasonable measure of the wage for new hires. Average posted wages in Burning Glass track quarterly measures of the average wage for new hires in survey data at the state, occupation, and industry level.

However, our dataset can measure not only averages but also *job-level* variation in posted wages—by which we mean the wage across successive vacancies posted by the same job title and establishment. Consider a physical location of Starbucks, in Cambridge, Massachusetts, that regularly posts vacancies for baristas and pays them an hourly wage. Our data tracks the hourly wage for baristas across multiple vacancies posted by the Starbucks. By studying job-level wages, we can purge wages changes due to job composition.

In the main contribution of the paper, we find that our measure of the wage for new hires is rigid downward, but flexible upward. We present two pieces of evidence. First, we detect signs of a constraint on wage setting. The posted wage rarely changes between successive vacancies of the same job—typically changing once every 5–6 quarters. When wages do change for a given job, they are four times more likely to rise than to fall. Many papers (e.g., McLaughlin 1994) document infrequent and asymmetric changes in continuing workers' wages; we show the same for a measure of new hires' wages.

Second, at the job level, the posted wage rises during expansions but does not fall during contractions. Figure 1 shows this result in a binned scatterplot. In the figure, the posted wage is averaged by job and quarter. On the vertical axis is wage growth between two consecutive vacancies for the same job. On the horizontal axis is the growth in quarterly state-level unemployment between the quarters in which the vacancies are posted.<sup>1</sup> As state unemployment decreases, the posted wage rises strongly, with wages responding similarly to small and large declines. As state unemployment increases, wages do not fall—neither for small nor for large increases. Figure 1 isolates job-level growth in posted wages to remove variation from changing job composition that might obscure downward wage rigidity. We confirm the finding with regressions and show wages are downwardly rigid with respect to identified labor demand shocks. Real wages are also rigid downward and

<sup>1</sup> Since many jobs do not post in consecutive quarters, sometimes the change in unemployment between postings is large. Average wage growth is greater than zero even during contractions, because some jobs experience positive wage growth during contractions.

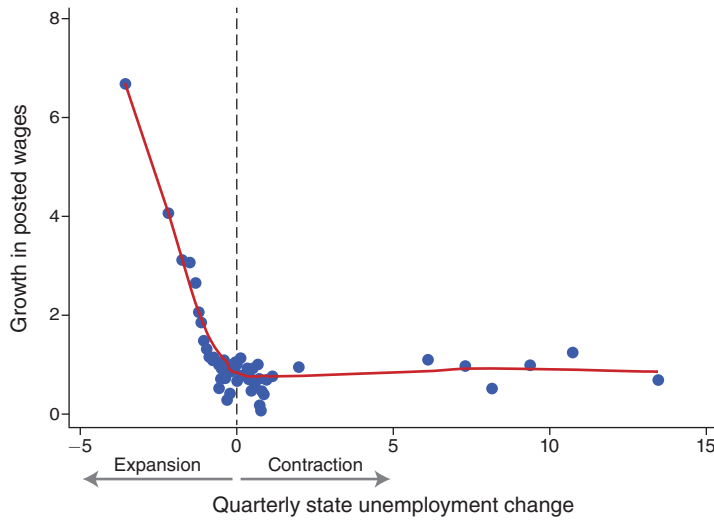


FIGURE 1. NOMINAL POSTED WAGE GROWTH AT THE JOB LEVEL AND UNEMPLOYMENT CHANGES

*Notes:* The graph plots wage growth of nominal posted wages, in percent, from Burning Glass; and state by quarter unemployment changes, in percentage points, from the local area unemployment statistics. The sample period is 2010:I–2020:II. To construct wage growth, we take the mean wage within each job and quarter, and then take log differences at the job level. We collect wage growth and unemployment changes into 100 bins, and add a nonparametric regression line.

flexible upward. If firms choose to post vacancies according to the business cycle, then our estimates might be subject to selection bias. However, we implement a standard Heckman (1979) selection correction, which suggests selection bias in the job-level regression is moderate.

Next, we show that job-level information is important for our results. The average wage for new hires, the object of some previous work, shows no sign of downward rigidity—neither using Burning Glass, nor worker-level survey data from prior work. We find that job composition raises the variance of average wages, meaning regressions using average wages lack the power to detect downward rigidity. Intuitively, average wages aggregate across all types of jobs. Then average wage changes reflect either wage changes at the job level, or changes in job composition. In the data, the share of low-wage jobs is volatile, which means average wages are also volatile. As a result, standard errors from regressions with average wages are twenty times larger than counterparts using job-level wages.

Our finding of downward wage rigidity for new hires is perhaps surprising. There is much evidence of downward wage rigidity for continuing workers—see, for example, Grigsby, Hurst, and Yildirmaz (2021) and Cajner et al. (2020) for convincing analyses during the Great Recession and the COVID-19 recession. However, wages could be more flexible for new hires than continuing workers. For example, firms might not cut wages for continuing workers to provide insurance, but still cut wages for new hires (Beaudry and DiNardo 1991). Our results instead suggest parity between new hires and continuing workers, perhaps due to internal equity (Bewley 2002).

Our findings can inform a range of models. For instance, wage rigidity for new hires is important in the canonical Diamond-Mortensen-Pissarides labor search model (e.g., Pissarides 2009). Nominal wage rigidity for new hires is also important in New Keynesian models with hiring (Basu and House 2016).<sup>2</sup>

There are four important limitations to our analysis. First, our main sample is small, being the subset of jobs in Burning Glass with wages, job title and establishment information, which post in multiple quarters; this sample is 0.8 percent of US vacancies. However, we show the main sample is broadly representative of the US population on observables, and wages from the main sample track wages from representative survey data. A second caveat is that we measure rigidity only for jobs that post wages. Jobs that do not post wages might be more flexible. Against this concern, we find the share of vacancies posting wages does not change during contractions. If vacancies that post wages were especially rigid, we might expect the share of vacancies posting wages to fall during recessions. A third caveat is that we do not measure the realized wage paid to new hires. Wage bargaining could lead to more flexibility in realized wages than in posted wages. However, we find wages are downwardly rigid even in occupations where wage bargaining is uncommon. A fourth caveat is that the COVID-19 recession is the main contraction over our sample period. Standard models focus on downward wage rigidity with respect to labor demand shocks, whereas the COVID-19 recession also involved labor supply shocks. However, in robustness we study identified labor demand shocks. We also end our main sample in June 2020, including the contractionary phase of the COVID-19 recession but not the aftermath, a period during which labor demand shocks were arguably important (Guerrieri et al. 2022).

*Related Literature.*—This paper contributes to three literatures. First, we contribute to the literature investigating the cause of unemployment fluctuations. Shimer (2005) shows that a standard calibration of the Diamond-Mortensen-Pissarides model leads to small unemployment fluctuations, compared with US data. Shimer (2004); Hall (2005); Hall and Milgrom (2008); Gertler and Trigari (2009); and Christiano, Eichenbaum, and Trabandt (2016) show that adding wage rigidity to the model leads to unemployment fluctuations as large as in the data. Pissarides (2009) emphasizes that in this model, the relevant wage is for newly hired workers. Our contribution is to argue wages for new hires are rigid downwards but flexible upwards.

This paper contributes to a second literature that measures wage rigidity for new hires. The seminal paper by Bils (1985) regresses the wage for new hires on unemployment to measure wage cyclicality. The wage data for new hires comes from surveys on workers switching jobs or entering new jobs from unemployment without job or establishment information. The regression controls for worker characteristics, but averages over the jobs into which workers are hired, which we term the average wage for new hires. Pissarides (2009) summarizes results from Bils (1985) and related papers. In these papers, point estimates suggest strong procyclicality, but the

<sup>2</sup>Rigidity in posted wages also matters in models of directed search (Moen 1997), since workers may form expectations from wage postings, which directs their search and affects vacancy creation. Continuing wages matter more than new hires' wages for unemployment fluctuations in other models, for instance with on the job search (Fukui 2020).

estimates tend to be imprecise and confidence intervals often include weak or zero procyclicality. Gertler and Trigari (2009) emphasize the challenge of job composition in interpreting these results.

Our work complements two papers that study wage rigidity for new hires and correct for job composition.<sup>3</sup> First, Gertler, Huckfeldt, and Trigari (2020) study wages for workers newly hired from unemployment. The average wage of workers hired from unemployment is likely less affected by job composition than the average wage of workers switching jobs. Gertler, Huckfeldt, and Trigari (2020) find weakly procyclical wages for workers hired from unemployment. Second, Grigsby, Hurst, and Yildirmaz (2021) introduce high quality payroll data on workers switching jobs. To control for job composition, they introduce a rich set of worker-level controls available in their data, and also develop an estimator that matches workers switching jobs to observationally identical workers at the destination firm. With either adjustment, the wage for new hires is weakly procyclical. Our paper complements these papers in two respects. First, our data is at the job level instead of the worker level, meaning we can directly correct for job composition. The worker-level controls in the two related papers may leave residual job composition. Second, both papers estimate that wages are rigid on average—we add evidence on the asymmetry of wage rigidity.

A third important paper is Martins, Solon, and Thomas (2012). Like us, the paper uses job-level wages, by exploiting rich administrative data from Portugal with establishment and detailed occupation information. The paper identifies “entry jobs” which frequently hire workers, and then measures the cyclicity of real wages paid to workers newly hired into those jobs. Like us, they find wages are similarly rigid for new hires and continuing workers, though all wages are more cyclical in Portugal than in the United States (see Carneiro, Guimarães, and Portugal 2012 for a related finding with Portuguese data).

Our paper contributes to a third literature that studies the consequence of downward rigidity for asymmetries in unemployment. For example, Dupraz, Nakamura, and Steinsson (2020) show that if wages for new hires are rigid downward and flexible upward, then unemployment rises sharply during contractions and falls more slowly during expansions. Our paper provides evidence for the asymmetric form of wage rigidity required by these papers, for a relevant kind of wage.<sup>4</sup>

## I. Data

This section introduces the main job-level dataset of posted wages. We study a dataset of wages posted on vacancies, with job title and establishment information. The dataset was developed by Burning Glass Technologies, a business analytics company. Burning Glass extracts information from the near-universe of job vacancy postings, using machine learning algorithms. There are 40,000 distinct online

<sup>3</sup>Other important papers studying the wage for new hires in US data include Haefke, Sonntag, and Van Rens (2013); Hagedorn and Manovskii (2013); Kudlyak (2014); Basu and House (2016); and Bils, Kudlyak, and Lins (2022).

<sup>4</sup>Acharya et al. (2018); Chodorow-Reich and Wieland (2020); Cacciatore and Ravenna (2021); Guerrieri et al. (2021); and Barnichon, Debortoli, and Matthes (2022) also study asymmetries in unemployment dynamics, when wages for newly hired workers are rigid downward and flexible upward.

sources, primarily job boards and company websites. No more than 5 percent come from any one source. The company employs a deduplication algorithm to avoid double counting vacancies that post on multiple job boards.

Our dataset covers January 2010 to June 2020. Therefore the dataset starts in the aftermath of the Great Recession, but excludes the Great Recession itself. Burning Glass began continuously collecting data in January 2010, meaning we cannot study earlier periods. Our dataset terminates after the contractionary phase of the COVID-19 recession, which officially took place over February–April 2020, but does not include the recovery after the contraction. We chose the end point to focus on a period in which labor demand shocks were relatively important, as we discuss at length in Section IIIB.

The dataset contains detailed information on the wage posted on vacancies, for around 20 percent of the full set of Burning Glass vacancies. Vacancies post either a wage range or a single wage. The data reports the pay frequency of the contract, for example, whether pay is annual or hourly; and the type of salary, e.g., base pay or bonus pay. Given pay frequency, we can measure hourly earnings for workers, i.e., the wage of the vacancy. The hours measure is an important advantage because in the United States, administrative data typically does not contain hours worked.

The data report establishment and job title information—henceforth, a “job” is a job title within an establishment whose wages are paid at a given frequency (e.g., annual or daily). These data are available for roughly half of the vacancies that post wages. Each physical location at which a firm employs workers is an establishment, measured by company name and zip code. Job titles are extracted from the text of the vacancies and cleaned using Burning Glass’ algorithms. The dataset also records any education requirements associated with the vacancy, such as high school diploma or undergraduate degree, if they are present. There is occupation information at the 2- 3- or 6-digit SOC code level, and the industry associated with the vacancy at the 2- or 3-digit NAICS level. Vacancies from 2018 onwards measure the length of time for which the vacancy was posted; and the source of the vacancy, such as a job board or company website.<sup>5</sup>

We now describe the steps toward forming our main job-level sample. The complete Burning Glass dataset has good coverage of the US economy, being the universe of online vacancies and 70 percent of total US vacancies, either online or offline (Carnevale, Jayasundera, and Repnikov 2014). However, we make significant restrictions that lead to a much smaller dataset, around 0.8 percent of total US vacancies.

Table 1, panel A, lists the sample restrictions, going from the complete Burning Glass dataset of online vacancies, in row 1, and ending with the main regression sample in row 6. Moving from the first to the second row, we restrict to the set of vacancies that post wage information, either single wages or wage ranges. This step reduces the sample size by 82 percent, as the number of observations in the second column indicates. Moving from row 2 to row 3, we exclude vacancies posting a wage range. The sample falls by an additional 63 percent. From row 3 to row 4, we drop vacancies that do not have job title and establishment information, reducing the

<sup>5</sup>See Supplemental Appendix Section C.1 for details of Burning Glass’ algorithms.

TABLE 1—SUMMARY STATISTICS

	Observations	Occupations	Counties
<i>Panel A. Sample formation</i>			
All Burning Glass	257,427,611	835	3,142
Range and point wages info	46,031,304	835	3,141
Point wage info	16,967,648	833	3,141
Point wage + job title	9,318,188	828	3,141
Point wage + job title, quarter collapse	8,990,542	828	3,141
Main sample	3,050,228	818	3,116
	Tenth percentile	Median	Ninetieth percentile
<i>Panel B. Cell counts for main sample</i>			
Observations per state	8,844	41,407	119,032
Observations per quarter	24,067	64,134	140,110
<i>Panel C. Information for main sample</i>			
Fraction	6-digit occupation	Length of posting	
Missing (percentage points)	0	76	
Fraction	Sector	3-digit industry	Vacancy source
Missing (percentage points)	14	22	67
Share by	Hourly	Annual	Monthly
Pay frequency (percentage points)	63	28	6
Share by	Base pay	Total pay	Bonus
Salary type (percentage points)	55	40	5
Length of	Tenth percentile	Median	Ninetieth percentile
Vacancy (days)	6	23	44
Nonzero	Tenth percentile	Median	Ninetieth percentile
Wage growth (percent)	-7.9	1.1	9.5

*Notes:* In panel A, an occupation is at the 6-digit SOC level, the definition of each part of the sample formation is described in the main text. In panel C, the first two rows report the share of vacancies missing information for each of the variables in the corresponding columns. We report only the three most common kinds of pay frequency and salary type. Wage growth is normalized by the number of quarters between the vacancy posting.

sample by an additional 45 percent. From row 4 to row 5, we take the mean wage within each job and quarter cell, which reduces the sample by 3.6 percent. From row 5 to row 6, we restrict to jobs that post in multiple quarters, leading to an additional 66 percent sample size reduction. Row 6 is our main sample, which we will study unless stated otherwise.

Table 1, panels B and C, also report cell counts and other information for the main regression sample. Panel B shows that there are a large number of observations for all states and quarters in the main regression sample. The main sample contains a mixture of hourly and annual jobs, and contains a mixture of base pay or total pay wages. Panel C shows that almost all of the data is classified into detailed occupations. However, there is more missing data for other fields, in part because Burning Glass collects several fields only after 2018. The length of vacancy posting, in the penultimate row of panel C, is similar to vacancy lengths from official sources.<sup>6</sup>

The restrictions to arrive at the main sample are necessary. To exploit the job-level information in our dataset, we must focus on vacancies with posted wages and job information. The jobs must be posted in multiple quarters, to allow a difference

<sup>6</sup>The duration of vacancies is similar to the mean vacancy duration reported in Davis, Faberman, and Haltiwanger (2013) from the BLS's JOLTS survey, which is 20 days.

within the job and over time. This sample is large in absolute terms, numbering some 3.05 million vacancies and covering almost all occupations and counties. However, the sample is small relative to the US population.

There are two concerns about measuring the wage for new hires using the main sample, which we try to address in Section IA. First, the main sample may not be representative of the US population, because jobs that repeatedly post wages or job title information may be different from others. Second, wages posted on vacancies may not be an accurate measure of the wage for new hires, perhaps because of ex post bargaining or stale information.

### *A. Posted Wages as a Measure of the Wage for New Hires*

We now show that our main sample of posted wages tracks measures of the wage for new hires from representative survey data—suggesting our main sample is reasonably representative and measures the wage for new hires adequately.

We construct an alternative measure of the wage for new hires from the Current Population Survey (CPS from CEPR 2025 and NBER 2025), for 2010:I–2020:IV, at the state, industry and occupation-level. The CPS is a worker-level survey that is representative of the US population. The wage for new hires is from workers switching jobs over the previous quarter, identified as in Fallick and Fleischman (2004); or workers entering jobs from unemployment, identified as in Haefke, Sontag, and Van Rens (2013). We take the mean wage at the state, 2-digit SOC occupation, or 2-digit NAICS industry level.

Then, we validate wages in the main Burning Glass sample at the state by quarter, industry by quarter, or occupation by quarter level, by regressing log CPS wages on log wages from our main sample. To avoid attenuation bias in the regression coefficients, due to measurement error in Burning Glass wages, we adapt the method of Angrist and Krueger (1995). We half the data in each state-quarter and calculate mean state-quarter wages in each subsample. When we take the mean, we reweight by 6-digit occupation shares in the 2016 Occupational Employment Statistics, in order to ensure that the Burning Glass data are representative of the population of occupations. We then instrument for wages in one subsample with the other. This procedure uncovers an unbiased estimate of the population coefficient from a regression of log CPS wages on log Burning Glass wages at the state level, provided that measurement error in wages is independently distributed. We adopt a similar procedure in our industry and occupation regressions.<sup>7</sup>

Table 2 shows that wages in the main Burning Glass sample follow CPS wages reasonably closely. Panel A shows state-level regressions, panel B shows industry-level regressions, and panel C shows occupation-level regressions. Column 1 is a simple bivariate regression; column 2 includes state fixed effects (in panel A), industry fixed effects (in panel B) or occupation fixed effects (in panel C); and column 3 includes time fixed effects. The regression coefficient in column 1 is 1.06. So, when Burning Glass wages change by one percentage, wages for new hires from the CPS also change by roughly one percent. Going across the columns, the point

<sup>7</sup>We include an additional two quarters of data of Burning Glass data, from 2020:III and 2020:IV, to compare Burning Glass and the CPS for the longest time period for which both datasets were available.

TABLE 2—AVERAGE NOMINAL WAGES IN BURNING GLASS VERSUS WAGE FOR NEW HIRES IN CPS

	Log average wage for new hires, Current Population Survey		
	(1)	(2)	(3)
<i>Panel A. Outcome is average wages by state</i>			
Log average wage, Burning Glass	1.06 (0.19)	1.09 (0.35)	0.85 (0.14)
Observations	2,244	2,244	2,244
Within $R^2$	0.06	0.03	0.05
<i>Panel B. Outcome is average wages by industry sector</i>			
Log average wage, Burning Glass	1.26 (0.20)	1.01 (0.23)	1.17 (0.23)
Observations	852	852	852
Within $R^2$	0.08	0.03	0.09
<i>Panel C. Outcome is average wages by broad occupation</i>			
Log average wage, Burning Glass	1.37 (0.15)	0.95 (0.18)	1.34 (0.15)
Observations	963	963	963
Within $R^2$	0.08	0.10	0.07
Time fixed effect			✓
Group fixed effect		✓	

*Notes:* This table presents regressions of average nominal wages in survey data on average wages in Burning Glass. The dependent variable is 100 times the log of the hours-weighted mean wage for newly hired workers from the 2010–2020 Current Population Survey, by quarter and state. We include both workers hired from unemployment (as in Haefke, Sontag, and Van Rens 2013) and workers hired from other jobs (as in Fallick and Fleischman 2004). Wages are trimmed at the first and ninety-ninth percentile. The wage is usual hourly earnings for hourly and nonhourly workers including overtime, for nonfarm workers; measured from the CPS Outgoing Rotation Group. In panel A, the dependent variable is the mean by state and quarter; in panel B, the dependent variable is the mean by quarter and 2-digit NAICS industry; in panel C, the dependent variable is the mean by quarter and 2-digit SOC occupation. The regressor is 100 times the log of average wages from 2010–2020 from Burning Glass. The sample is the sample of Table 4, i.e., the main regression sample, with additional quarters of data 2020Q3 and 2020Q4 from Burning Glass. Salaries are trimmed at the fifth and ninetieth percentiles in each 2-digit occupation group, within each pay frequency and salary type. To uncover group-quarter salaries, we regress  $\log(\text{salary}_{ist}) = \alpha + \sum_{p,s} \beta_{ps} D_{ps} + \sum_{s,t} \gamma_{st} W_{st} + \text{error}_{ist}$  where  $D_{ps}$  denotes a set of salary type by pay frequency dummies and  $W_{st}$  is a set of group (i.e., state, 2-digit industry or 2-digit occupation) by quarter dummies. Observations are reweighted by the 2014–2016 OES (see Occupational Employment Statistics 2025). Then  $W_{st}$  is the log mean salary in the group-quarter. We split the sample in half in each state-quarter, and instrument for salaries in one subsample with salaries in the other, to overcome measurement error. The regressions are weighted by the number of CPS observations in each group (i.e., state, industry, or occupation) and quarter, using sampling weights. Standard errors are clustered by group and quarter.

estimate is not systematically different and remains close to 1 as we add in time or group fixed effects. Going down the panels, the fit is similar at the state, industry or occupation level. We cannot reject the coefficient equals 1 in most specifications.

Comparing average wages in the CPS and Burning Glass is useful, even though the CPS measures wages at the worker level, and Burning Glass at the job level. The reason is that *average* wages for new hires, in both job-level and worker-level

datasets, should measure the same object. After all, all newly hired workers must be hired into a newly created job. Therefore the average wage for new hires, aggregating across workers, should equal the average wage for new hires, aggregating across jobs. If average wages in Burning Glass did not closely comove with average wages in the CPS, we might suspect mismeasurement in Burning Glass.

Still, this exercise has limitations. The partial  $R$ -squared of the regression is sometimes high, for example 0.1 in the broad occupation regression with group fixed effects; but sometimes low, for example 0.03 in the state regression with state fixed effects; and the standard errors are relatively large in column 2 of panel A. These patterns are not surprising, because survey wages are known to be noisy.<sup>8</sup>

There are likely three reasons why wages attached to online vacancies seem to measure the wage for new hires. First, for a representative survey of job-seekers, Hall and Krueger (2012) report that at least 30 percent and as many as 65 percent of workers do not bargain over the wage of the new vacancies to which they apply, and instead receive a wage dictated to them by their employer when they are hired.<sup>9</sup> Therefore for many newly hired workers, the wage attached to the vacancy is the relevant wage at the start of the match. Second, online vacancy posting is costly, which discourages firms from posting out-of-date wage information. The median cost of posting a vacancy on the largest four online job boards was \$419 in 2017, and companies posting on their own websites typically pay monthly fees to subcontractors.<sup>10</sup> Gavazza, Mongey, and Violante (2018) show that company websites and online job boards are a large share of total recruiting costs for the typical US firm. Third, the duration of vacancies is short, which prevents “stale” vacancies. The median time for which vacancies are posted is 23 days, and 92 percent of vacancies are removed within the quarter.

To further address concerns about representativeness, Supplemental Appendix Section C.2 provides a detailed comparison between our main sample and official, population-level data. To summarize, our main sample seems reasonably representative of the population. The occupation and industry mix as well as establishment characteristics are similar, and selection into the main sample is not correlated with business cycles. The main caveat is that the main sample under weights healthcare and over weights transportation and retail. In robustness exercises for our empirics, we reweight to the occupational or regional distribution of jobs, and find little change to our results.

### B. Advantages of Job Level Data

Our job-level dataset has an advantage relative to worker-level survey data measuring the wage for new hires. We can track a measure of the wage for new hires at

<sup>8</sup>A concern is that the state-level regression with time fixed effects has a coefficient of 0.85, well below 1. We suspect that this gap is sampling error, given that the other regressions with time fixed effects are at or above 1. Supplemental Appendix Table 1 regresses the difference between state-level wages in Burning Glass and the CPS on state unemployment and time fixed effects, and finds that the difference between these wages is uncorrelated with unemployment.

<sup>9</sup>Hall and Krueger (2012) find that 30 percent of workers knew their exact wage before being hired, and above 80 percent of workers “knew exactly or had a pretty good idea” of their wage before being hired.

<sup>10</sup>Kehayas, Cailen. 2018. “How Much It Costs to Post A Job Online (The Dirty Truth).” <https://web.archive.org/web/20220119141428/https://blog.proven.com/how-much-to-post-a-job>.

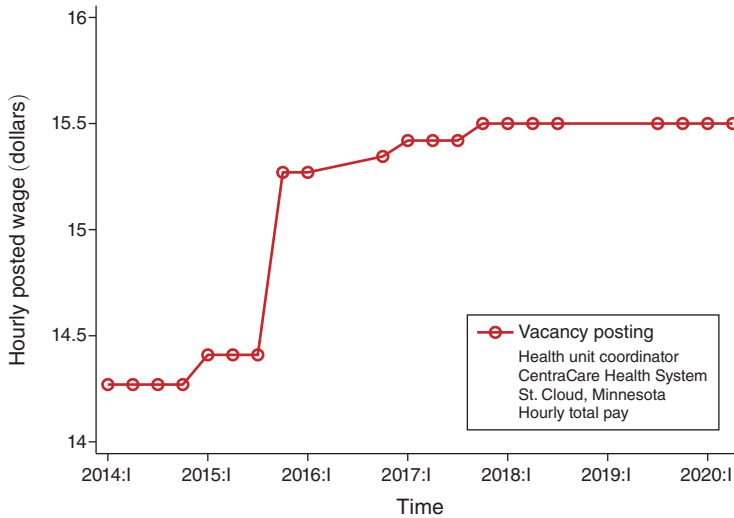


FIGURE 2. AN EXAMPLE OF A JOB

Notes: This graph plots an example of a job-level wage. The hourly nominal posted wage is the mean within each quarter in which the job posts a vacancy.

the job and the establishment level. We can track wages across multiple vacancies posted by the same job, within the same establishment. In coming sections, we use this feature to document downward rigidity in a measure of the wage for new hires faced by establishments.

Figure 2 displays job-level variation. We present a job that posts multiple vacancies. The firm is CentraCare Health System. The establishment is the branch of the firm in St. Cloud, Minnesota. The job title is Health Unit Coordinator. The salary is hourly total pay. When the vacancy posts multiple times within the quarter, we take the average. Then according to our definition, a job is a health unit coordinator at the St. Cloud establishment of CentraCare. The job posts vacancies in 21 quarters over seven years. We can track the wage across these vacancies; that is, we can track job-level changes in the posted wage.<sup>11</sup>

Worker-level data cannot easily track job-level variation in the wage for new hires. Workers are typically hired once into a job. So worker data cannot easily track the wage across successive workers, hired into the same type of job. Survey measures of wages, such as the Current Population Survey or the National Longitudinal Survey of Youth, typically measure workers' wages and do not contain job or establishment information.

In the sections to come, the job-level data will let us document new findings about wage rigidity for new hires. However, job-level data has two further and related benefits. First, job-level data can control for wage changes due to job composition. We will see that measures of the wage that do not adjust for composition behave

<sup>11</sup>More generally, our measure of a job seems to capture the relevant heterogeneity in wages due to job-level characteristics. In Supplemental Appendix Table 2, we regress log wages on job fixed effects, finding that job fixed effects explain at least 82 percent of the variation in wages, meaning that the tasks associated with a given job title and establishment are relatively stable over time.

differently. Second, in standard models, variation in the wage for new hires at the job level is key for unemployment fluctuations. Therefore we are measuring a particularly important object in the data. This point is well known in the labor search literature (e.g., Gertler, Huckfeldt, and Trigari 2020).

## II. Constraint on Wage Setting for New Hires

This section presents evidence of a constraint on wage setting. We document that posted wages typically do not change between successive vacancies of the same job. Moreover when wages do change for a given job, they rarely fall.

### A. Hazard Estimation of the Probability of Wage Changes

For the job in Figure 2, there is a distinctive pattern: The wage changes infrequently across vacancies. When the wage does change, it always rises. We now provide more systematic evidence of this pattern.

Measuring the frequency of wage change is difficult because the wages of a job are “missing” in the quarters without vacancy posting. Therefore, we cannot directly observe the probability that the wage for new hires changes. We adapt a standard approach from the price setting literature to overcome this problem (Nakamura and Steinsson 2008; Klenow and Kryvtsov 2008). We treat the wage as a latent variable which evolves stochastically when it is unobserved, and treat the observed sequence of wages as draws from the latent process. We estimate the latent process with a constant hazard model. We can then calculate the probability that the wage changes, even if jobs do not post in all quarters.

The constant hazard model has several desirable properties. If the observed wage does not change between successive vacancies, the latent wage also does not change. If the observed wage does change, the latent wage also changes. The latent wage can change multiple times if the observed wage changes once, and is more likely to change if the gap between successive vacancies is longer.<sup>12</sup> One can easily adapt this process to separately estimate the probability of wage increase and decrease. We use implied durations to measure for how long wages are unchanged. Following Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008), we view the constant hazard model as a convenient way to summarize the data. However, the model will not accurately represent all features of the data generating process if, for example, there is state dependence in wage setting.

### B. Infrequent Wage Changes at the Job Level

We find that the nominal posted wage changes infrequently, suggesting a constraint on wage setting for new hires at the job level. Table 3 reports the results. Row 1 estimates the quarterly probability of wage change according to our method. Row 2 reweights vacancies at a granular level, to target the distribution of jobs from the 2016 Occupational Employment Statistics. Row 3 reweights to target the state

<sup>12</sup>Supplemental Appendix Section C.3 describes the hazard estimation procedure in full.

TABLE 3—QUARTERLY PROBABILITY OF NOMINAL POSTED WAGE CHANGE AT THE JOB LEVEL

	Prob. change (1)	Duration of unchanged wages (2)	Prob. decrease (3)	Prob. increase (4)	Observations (5)
Unweighted	0.17	5.36	0.04	0.12	1,746,711
Occupation weight	0.19	4.79	0.04	0.13	1,746,711
Region weight	0.17	5.45	0.04	0.12	1,746,711
No low-wage jobs	0.18	5.07	0.03	0.12	1,266,461

*Notes:* We study the main sample of Burning Glass data. We estimate the probability of job-level wage change using a similar method to Nakamura and Steinsson (2008). We assume that the hazard rate of job change/increase/decrease is constant and identical for all jobs in the same 2-digit SOC code occupation. We then estimate the hazard rate of job change/increase/decrease by maximum likelihood. We then calculate the implied duration and probability of change/increase/decrease for each occupation, and then take the median across occupations, weighted by the number of vacancies. In row 2, we reweight to target the distribution of jobs at the 6-digit SOC level from the 2014–2016 OES. In row 3 we reweight to target the distribution of employment across states from the 2010 QCEW (see Quarterly Census of Employment and Wages 2025). In row 4 we drop jobs in the bottom quartile of the wage distribution. Counts refer to the number of differenced observations.

distribution of jobs from the Quarterly Census of Employment and Wages. Row 4 drops jobs from the bottom quartile of the wage distribution, since minimum wages might cause infrequent changes. Across all columns, the probability of wage change is similar and low, between 0.17 and 0.19 per quarter. The corresponding implied durations of unchanged wages are four to five quarters.

The posted wage changes infrequently even as many successive vacancies are posted for a given job. Supplemental Appendix Table 3 reports, for each vacancy, the length of time that has elapsed since a previous vacancy was posted for the same job. For over 90 percent of vacancies, less than five quarters have elapsed since the job posted a previous vacancy. Therefore jobs often post vacancies several times, over multiple quarters, without changing the wage. Data tracking individual workers’ wages cannot easily measure the frequency of wage change for new hires. Workers are typically hired once into a job, but the object of interest is the wage across successive workers hired into the same type of job.

Infrequent changes in posted wages suggest a constraint on wage setting. We now present evidence suggesting this constraint is asymmetric. Figure 3 plots the distribution of wage growth, after removing vacancies with zero wage change. We take the distribution of wage growth for new hires between two consecutive vacancies posted for the same job, and then exclude observations with zero wage growth. As before, we average wages within each job-quarter, meaning wage growth is quarterly. We truncate the plot at ±20 percent wage growth.

There are two clear points. First, conditional on a wage changing, wages for new hires rise more often than they fall. Secondly, wages “pile up” close to zero: There are many small positive wage increases, but far fewer small wage decreases. Both points suggest a downward constraint on wage setting for new vacancies of a given job.

We then estimate the probability of wage increases and decreases for new hires. The results are in Table 3, columns 3 and 4. As expected, wages are more likely to rise than to fall. For example, in the unweighted specification of row 1, the probability of wage decreases is 0.04, whereas the probability of wage increases is 0.12.

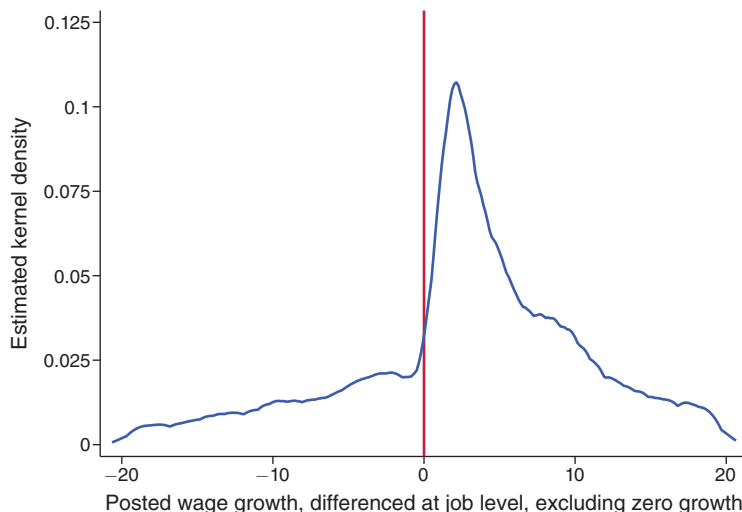


FIGURE 3. DISTRIBUTION OF NOMINAL POSTED WAGE GROWTH, NONZERO WAGE GROWTH ONLY

*Notes:* This graph is the distribution in the growth of nominal posted wages, excluding zeros, from Burning Glass, with a kernel density. Wages are averaged by job and quarter. Wage growth is the growth in wages between two consecutive vacancies posted by the same job, measured in percentage points. The wage growth distribution is truncated at  $\pm 20$  percent. The sample period is 2010:I–2020:II.

Table 3 shows that the finding is similar across several specifications, including after reweighting to target the occupational or regional distribution of jobs, or excluding low-wage jobs to remove the effect of minimum wages.

Overall, the pattern of infrequent changes and especially rare decreases suggests an asymmetric constraint on wage setting. An alternative explanation for these patterns relates to vacancy fill rates. Vacancies posting a high wage might be filled quickly, whereas vacancies posting a low wage might be filled more slowly. Therefore, jobs that initially post vacancies with a low wage might raise wages in the future, in order to raise vacancy filling rates. As a result, wage increases might be more likely than wage decreases, regardless of any constraint.

We can directly test for the relevance of vacancy fill rates using information on the length of time for which vacancies are posted. This variable, which is collected by Burning Glass, is a proxy for vacancy fill rates. Supplemental Appendix Table 4 finds that the probability of wage increase and decrease is similar for vacancies that are posted for long or for short periods of time. This result suggests behavior relating to vacancy fill rates cannot explain the pattern of wage increases and decreases that we estimate.

Our finding, that a measure of the wage for new hires changes infrequently and falls especially rarely, has not been previously documented. We provide context with a better known fact: Workers in continuing employment rarely experience wage changes.

Figure 4 shows that the duration of unchanged wages is similar for continuing jobs and for posted wages. The figure presents estimates from four papers that are representative of the literature, of the length of time for which wages are unchanged in continuing jobs. The first two papers, Card and Hyslop (1997) and Daly and Hobbijn

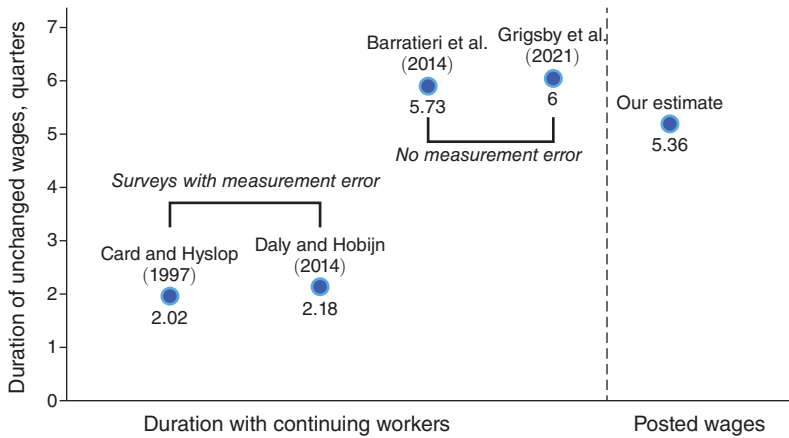


FIGURE 4. DURATION OF UNCHANGED NOMINAL WAGE FOR CONTINUING WORKERS AND POSTED WAGES

Notes: This graph plots the implied duration for which wages are unchanged from four papers that study continuing wages using payroll and survey data, alongside our estimate for posted wages using Burning Glass data.

(2014), construct the frequency of wage change from worker-level survey data, which is thought to overestimate the frequency of wage change due to measurement error. The third estimate, from Barratieri et al. (2014), studies worker-level data with a correction for measurement error. The final paper, Grigsby, Hurst, and Yildirmaz (2021), studies high quality payroll data. Two estimates are close to ours: the estimates of Grigsby, Hurst, and Yildirmaz (2021) and Barratieri et al. (2014), which are both unaffected by measurement error.

The probability of wage cuts is also similar for posted versus continuing wages. Over 2008–2016, Grigsby, Hurst, and Yildirmaz (2021) report that 5 percent of base wage changes are wage cuts, and around 25 percent of changes in total wages (including bonus and overtime) are wage cuts. For a similar period, Kurmann and McEntarfer (2019) study high quality payroll data from the state of Washington, and find that around 25 percent of changes in total wages are wage cuts. Over the COVID-19 recession, Cajner et al. (2020) report that 13 percent of base wage changes are wage cuts. We find a number in this range: That is, 23 percent of wage changes are wage cuts.<sup>13</sup>

Our findings suggest that new and continuing wage changes are broadly governed by the same considerations which relates to a debate on how wages are set for new hires. Some previous work conjectures parity in wage setting due to internal equity between new hires and continuing workers (Bewley 2002). Our finding for posted wages supports this argument. However, other plausible mechanisms predict differences between new hires and continuing workers. As one example, implicit contracting models imply that continuing wages should be rigid downwards, while wages for new hires should be *flexible* downwards (Beaudry and DiNardo 1991).

<sup>13</sup>The summary statistics in Table 1 show that our data contains a mixture of base and total wages.

As a second example, continuing workers might have a reference point of their own past wage, and object to wage cuts because of morale. These considerations might matter less for new hires (Eliaz and Spiegler 2014).

### III. Wage Cyclicity

Motivated by our previous finding of asymmetric wage changes, this section asks whether posted wages respond differently to contractions versus expansions in the business cycle. We find that, across successive vacancies posted by the same job, the nominal posted wage does not fall during contractions, but does rise during expansions.

We study the response of wages to unemployment, following Bilts (1985), and focus on regional business cycles. There were three major regional business cycles over 2010–2021. First, there was the COVID-19 recession, which had an uneven regional impact (Chetty et al. 2020). Second, there was a boom and bust in natural resource producing states, such as Texas, due to fluctuations in the global oil price between 2010 and 2016. Third, our sample includes the recovery from the Great Recession, though the sample excludes the Great Recession itself.

We estimate a variant of the standard regression of Bilts (1985) and add an additional term for asymmetry, motivated by our previous findings. We estimate the cyclicity of posted wages at the job level as

$$(1) \quad \Delta_{t,t-j} \log w_{ist} = \alpha + \gamma_t + \zeta_t I[\Delta_{t,t-j} U_{st} < 0] + \beta \Delta_{t,t-j} U_{st} \\ + \delta I[\Delta_{t,t-j} U_{st} < 0] \Delta_{t,t-j} U_{st} + \varepsilon_{is,t,t-j}.$$

Here,  $w_{ist}$  is the nominal posted wage in job  $i$  and quarter  $t$ . Using the difference operator  $\Delta_{t,t-j}$ , we difference wages between two consecutive quarters  $t-j$  and  $t$  in which the job posts a vacancy—this gap may be more than a single quarter. This step isolates job-level wage changes.  $\Delta_{t,t-j} U_{st}$  is the change in quarterly state unemployment, differenced over the same quarters as the vacancies, and  $I[\Delta_{t,t-j} U_{st} < 0]$  is an indicator variable for whether unemployment is falling.  $\gamma_t$  is a time fixed effect, which isolates regional variation; we also interact the time fixed effect with an indicator for whether unemployment is rising or falling, denoted by  $\zeta_t I[\Delta_{t,t-j} U_{st} < 0]$ .  $\beta$  and  $\delta$  measure the sensitivity of the posted wage to regional unemployment. A more negative number indicates greater sensitivity. If  $\delta < 0$ , then wages comove more with unemployment during expansions than contractions. If  $\beta = 0$ , then wages do not comove with unemployment during contractions.

#### A. Asymmetric Response of Wages to Unemployment

We now introduce the key empirical result. Figure 1, previously shown in the introduction, shows that when unemployment rises, the nominal posted wage does not fall—whereas wages do rise as unemployment falls. Moreover, the comovement of wage growth and unemployment changes is piece-wise linear around zero unemployment changes. Therefore, nominal wages respond similarly to large and small changes in unemployment.

TABLE 4—NOMINAL POSTED WAGE GROWTH AND UNEMPLOYMENT CHANGES AT THE JOB LEVEL

	Nominal wage growth at the job level, $\Delta \log w_{jst}$			
	(1)	(2)	(3)	(4)
$\Delta U_{st}$	0.05 (0.01)	-0.05 (0.03)	-0.05 (0.03)	-0.65 (0.10)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.40 (0.10)	-1.68 (0.11)	-1.68 (0.10)	
Observations	1,789,042	1,789,042	1,789,042	1,789,042
Time fixed effect		✓	✓	✓
State fixed effect			✓	

*Notes:* This table presents estimates regressing nominal job-level wages on unemployment. The dependent variable is quarterly percent growth in nominal posted wages, from the Burning Glass main sample. Wage growth is trimmed at the first and ninety-ninth percentiles. The regressors are the change in state-quarter unemployment, and the change interacted with an indicator for whether unemployment is decreasing, from the 2010–2020; see Local Area Unemployment Statistics (LAUS 2025), in percentage points. Standard errors are in parentheses, clustered by state. Column 1 presents estimates without controls. Column 2 adds in time fixed effects, interacted with an indicator variable for whether unemployment is decreasing. Column 3 adds in state fixed effects. Column 4 presents estimates without asymmetries, by only including unemployment changes and time fixed effects as regressors. Counts refer to the number of differenced observations.

Table 4 confirms these results by estimating regression (1). Column 1 of the table estimates the regression without time fixed effects. The coefficient on  $\Delta U_{st}$  has a slightly positive value of 0.05—meaning the posted wage does not respond to contractions. The coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  has a significantly negative value of -1.4. Therefore, wages are more sensitive to expansions than contractions in unemployment. Summing the coefficients shows that wages are flexible upward: After a fall in unemployment of 1 percentage point, wages grow by 1.4 percent. The results are robust across several specifications. Column 2 is our preferred specification, which adds time fixed effects in order to isolate regional variation. The coefficient on  $\Delta U_{st}$  becomes slightly negative, with a value of -0.05, and the coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  is even more negative than in column 1. Column 3 adds state-specific trends and finds a similar result. Column 4 drops the  $I[\Delta U_{st} < 0] \Delta U_{st}$  term from our benchmark regression to estimate the average sensitivity of wage growth to unemployment changes, while still including time fixed effects. On average wages do comove negatively and significantly with unemployment. The coefficient of -0.65 implies that a 1 percentage point fall in unemployment raises wages by 0.65 percent. However, this average comovement is entirely driven by expansions and not contractions.

So far, we have studied nominal wages. Real wages are more relevant in many theoretical models and could behave differently from nominal wages depending on price inflation. We now show there is similar rigidity and asymmetry in real wages. We study a version of regression equation (1), but replace the outcome with real wages. Real wages are deflated using regional prices measured at the census division level. We study two measures of the consumer price index, either including or excluding shelter (US Bureau of Labor Statistics 2025). Correspondingly, we measure unemployment at the census division level for the analysis with real wages.

TABLE 5—REGRESSION OF POSTED WAGE GROWTH ON CENSUS DIVISION UNEMPLOYMENT CHANGES

	Wage growth at the job level		
	Nominal wages	Real wages (price level excludes shelter)	Real wages (price level includes shelter)
$\Delta U_{st}$	-0.13 (0.04)	-0.15 (0.09)	-0.10 (0.08)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.85 (0.14)	-1.02 (0.05)	-0.37 (0.14)
Observations	1,696,741	1,696,560	1,696,552
Interacted time fixed effect	✓	✓	✓

*Notes:* In all panels, the dependent variable is quarterly percent growth in posted wages, from the Burning Glass main sample. Wages are averaged within each job and quarter. Wage growth is trimmed at the first and ninety-ninth percentiles. In column 1, the outcome is nominal wages. In column 2 the outcome is real wages, deflated by the census division level consumer price index (see Supplemental Appendix Section C.4 for details on how we construct price indices). In column 3 the outcome is real wages, deflated by the census division level consumer price index excluding shelter. The regressors are the change in census division-quarter unemployment, in percentage points, and an interaction of the change with an indicator for whether unemployment is falling, from the 2010–2020 LAUS, in percentage points. Standard errors are in parentheses, clustered by census division. We control for time fixed effects, interacted with an indicator variable for whether unemployment is decreasing. Counts refer to the number of differenced observations.

Using regressions, Table 5 shows that real wages are rigid downward and flexible upward. We estimate our preferred specification with time fixed effects. For comparison, column 1 of the table reestimates nominal wage cyclicality at the census division level. Nominal wage rigidity is similar to the previous state-level estimates. Column 2 then estimates real wage cyclicality, deflating by consumer prices excluding shelter. The coefficient on  $\Delta U_{st}$  of  $-0.15$  implies that when unemployment rises by 1 percentage point, real wages fall slightly. However, the coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  is significantly negative, meaning real wages are more sensitive to expansions than to contractions. Summing the coefficients shows that real wages are flexible upward: In response to a fall in unemployment of 1 percentage point, real wages grow by 1.02 percent. Column 3 estimates real wage cyclicality deflating by all consumer prices including shelter. The degree of downward wage rigidity is similar to column 2: The coefficient on  $\Delta U_{st}$  changes slightly to  $-0.10$ . The degree of upward wage flexibility is smaller, though still significantly negative, since the coefficient in column 3 is closer to zero with a value of  $-0.37$ .

The coefficient on  $\Delta U_{st}$  is similar in all three specifications, meaning nominal and real wages are both downwardly rigid. The coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  is closer to zero for real wages than for nominal wages. However, real wages are still flexible upwards, especially if the price level does not include shelter.<sup>14</sup> The similar behavior of nominal and real wages is probably because regional prices are sticky (Hazell et al. 2022).

<sup>14</sup> Measures of real wages that exclude shelter are likely more relevant for labor search models, since the relevant real wage depends on producer prices, which do not include housing.

### B. Wage Rigidity with Respect to Labor Demand Shocks

We now turn to an identification challenge. In labor search models, the object of interest is typically wage rigidity with respect to labor demand shocks. However, the major contraction during our sample period is the COVID-19 recession, which also included labor supply shocks—for instance from fears of disease. Labor supply shocks could mean that wages do not fall during contractions, even if there were no constraint on wage setting. We now describe various strategies that attempt to isolate labor demand variation.

Our first strategy to include labor demand variation in our regressions is to end the main sample in June 2020. Therefore the sample includes the contractionary phase of the COVID-19 recession, which officially took place during February–April 2020, but does not include the aftermath of the contraction. As such, we try to include the portion of the pandemic period during which labor demand shocks were arguably important. Supplemental Appendix Figure 1 plots the employment-population ratio from the start of 2019 onwards. The first half of 2020 contains a sharp contraction and rebound, whereas the period after June 2020 is a more gradual recovery. Existing theoretical and empirical work suggests that a large part of the sharp contraction in employment was due to labor demand. For instance, Guerrieri et al. (2022) study a two-sector model with downward nominal wage rigidity and incomplete markets. They argue that under plausible conditions, a shutdown of one sector represents a contractionary labor demand shock to the sector that remains open. Baqaee and Farhi (2022) consider a model with multiple sectors, input-output linkages, and downward nominal wage rigidity. They find that labor demand shocks can explain half of the contraction in real GDP during the COVID-19 recession. Empirically, Forsythe et al. (2020) show that declines in vacancy posting were similarly large for in-person versus work-from-home occupations—the latter jobs are less likely to be affected by labor supply shocks.

Supplemental Appendix Figure 2 presents suggestive evidence that labor supply shocks may have become more important after June 2020. This figure repeats the main figure of the paper, Figure 1, in blue circles. In red triangles, we add observations from July 2020–March 2021, which were also available at the time of writing. The data from July 2020 onwards suggests supply shocks, because wage growth is higher on average, and uncorrelated or even positively correlated with unemployment changes.

We now use four strategies to isolate labor demand shocks before June 2020, and tentatively argue there is downward wage rigidity with respect to these shocks. Our first strategy adds controls for labor supply shocks to the baseline regression equation (1). The controls are designed to absorb observable supply shocks, such as the effect of unemployment insurance (UI) extensions under the CARES act. However, this strategy will not work well if we have failed to control for a large labor supply shock.<sup>15</sup>

<sup>15</sup>Specifically, we add as controls, the interaction of an indicator variable for 2020 onwards, with: the length of state shutdowns; quarterly COVID-19 case and death counts; the state's median UI replacement after the CARES act; and the state's 2010–2019 share of female employment.

The second strategy instruments for unemployment with a standard industry “shift share” measure of labor demand. The instrument is a weighted average of national industry employment growth, with weights depending on the state’s employment share in each industry. The identification assumption is that differences in national employment growth across industries are due to labor demand and not labor supply shocks. Supplemental Appendix Table 5 estimates the first stage regression corresponding to this instrument, and finds a strong first stage both for contractions and for expansions as measured by the instrument. As a by product, this instrument also corrects attenuation bias due to measurement error in state unemployment.<sup>16</sup>

The third strategy classifies sectors or occupations that, according to theory, were affected by labor demand and not labor supply shocks. We use two different classifications. The first classification follows Guerrieri et al. (2022), who suggest sectors that were not shut down were subject to labor demand shocks. Therefore, we reestimate our baseline regression only for occupations amenable to remote work (Dingel and Neiman 2020). The second classification follows Baqaee and Farhi (2022) who identify a set of sectors that faced labor demand shocks, in a calibration of their model. Importantly, their calibration strategy does not use information on wages.

Our final strategy develops another shift share instrument based on regional exposure to oil shocks, which uses variation only from before the COVID-19 recession. Specifically, we instrument for unemployment with the quarterly global oil price interacted with state specific indicators and indicators for whether the oil price was increasing or decreasing. This regression contains many instruments and therefore is estimated by limited information maximum likelihood (LIML) as Angrist and Frandsen (2022) recommend. This regression only includes data from before 2020. This instrument exploits the large variation in the global oil price before 2020—Supplemental Appendix Figure 4 shows that the oil price increased substantially between 2010–2014 and then fell sharply. The identification assumption is that oil price shocks particularly affect labor demand in states concentrating in natural resources, such as Texas (Acemoglu, Finkelstein, and Notowidigdo 2013).

Table 6 shows the results. Reassuringly, across a range of identification strategies, wages seem to be rigid downward and flexible upward. The first column is the baseline regression estimates, which do not attempt to identify labor demand. Column 2 adds labor supply controls to the regression. Both regression coefficients change by a small amount. Column 3 studies the industry shift share instrument. The coefficient on  $\Delta U_{st}$  declines slightly to  $-0.45$ , though is less precisely estimated. The coefficient is not significantly different from zero and still suggests downward rigidity, albeit less than the baseline estimate. With the shift share instrument, the degree of asymmetry is even greater, since the coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  takes a more negative value of  $-2.47$ . Columns 4 and 5 estimate wage cyclicality only for work from home occupations and for industries facing demand shocks, respectively. In both cases the degree of wage rigidity is

<sup>16</sup>Supplemental Appendix Figure 3 suggests variation in industry employment growth during the COVID-19 recession was primarily determined by labor demand. For example, gambling had large employment falls, whereas home building materials had a large rises, which suggests changes in labor demand from the goods market. See Chodorow-Reich et al. (2022) for a similar instrument.

TABLE 6—NOMINAL POSTED WAGES AND UNEMPLOYMENT: IDENTIFYING LABOR DEMAND VARIATION

	Nominal wage growth at the job level, $\Delta \log w_{jst}$						
	Baseline	Labor supply controls	Shift share instrument	Work from home only	Demand shock industries	Oil instrument (pre-2020)	Baseline (pre-2020)
$\Delta U_{st}$	-0.05 (0.03)	-0.08 (0.03)	-0.45 (0.24)	-0.05 (0.03)	-0.02 (0.04)	1.45 (1.77)	-0.32 (0.12)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.68 (0.11)	-1.64 (0.10)	-2.47 (0.32)	-1.35 (0.13)	-1.72 (0.12)	-3.12 (1.84)	-1.38 (0.13)
Observations	1,789,042	1,781,994	1,789,042	442,895	323,205	1,535,009	1,535,009
Robust F stat.	65.4						
Time fixed effect	✓	✓	✓	✓	✓	✓	✓
Estimation method	OLS	OLS	2SLS	OLS	OLS	LIML	OLS

Notes: This table presents estimates regressing nominal job-level wages on unemployment, while identifying labor demand variation. Column 1 presents the baseline estimate from column 2 of Table 4. Column 2 repeats column 1 but controls for a dummy variable for 2020:I onwards, interacted with: the length of time for which states shut down businesses (from Chetty et al. 2020); quarterly COVID-19 case and death counts (from Chetty et al. 2020); the state’s median UI replacement after the passage of the CARES act (from Ganong, Noel, and Vavra 2020); and the state’s 2010–2019 share of female employment (from the Current Population Survey). Column 3 instruments unemployment changes with a shift share instrument, and the instrument interacted with an indicator for whether it is increasing; and also controls for time fixed effects interacted with whether the shift share instrument is increasing; for 2010–2020. The shift share instrument is measured at the 6-digit NAICS level from the QCEW, we calculate state industry shares using annual employment from 2009. Column 4 reestimates the regression of column 1, but restricts to work-from-home occupations as classified by Dingel and Neiman (2020). Column 5 reestimates the regression of column 1, but restricts to occupations affected by labor demand shocks as classified by Baqaee and Farhi (2022). Column 6 instruments unemployment changes with an instrument based on the growth in the Brent Crude oil price, averaged by quarter, and controls for time fixed effects interacted with whether the oil price is increasing. The instrument is a set of variables  $A_s \Delta \log \text{oil price}_t + B_s \Delta \log \text{oil price}_t \times I(\Delta \log \text{oil price}_t < 0)$ , where  $A_s$  and  $B_s$  are state fixed effects estimated in the first stage. In this column, the sample is 2010–2019. Column 7 is the same as column 1 but restricts the sample to 2010–2019. Each regression is estimated by OLS; with the exception of the shift share instrument (estimated by 2SLS) and the oil instrument (estimated by LIML). Counts refer to the number of differenced observations.

similar to the baseline. Column 6 estimates wage rigidity instrumenting for unemployment with the oil price instrument, using only variation from before 2020. The estimates are less precise but broadly agree with the others—the coefficient on  $\Delta U_{st}$  is positive, which suggests that wages do not fall in response to contractions. Again, there is a large negative coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$ . Column 7 reestimates the baseline specification excluding data from 2020 onwards, therefore restricting to a period where labor supply shocks were likely less important. The coefficients again suggest that wages were rigid downward and flexible upward, with similar magnitudes to the baseline.

### C. Downward Wage Rigidity: Robustness Tests

The remainder of this section studies various robustness tests. Table 7 groups together some robustness tests about our key finding. Each row estimates versions of our benchmark regression, reporting the coefficient on  $I[\Delta U_{st} < 0] \Delta U_{st}$  and its standard error. If this coefficient is negative, then wages are more rigid downward than upward.

TABLE 7—ROBUSTNESS—NOMINAL POSTED WAGES ON UNEMPLOYMENT AT THE JOB LEVEL

Specification	Coefficient $\Delta U_{st} \times I(\Delta U_{st} < 0)$	Standard error	Observations
1. Baseline	-1.68	0.11	1,789,042
2. Annual	-1.58	0.13	809,187
3. Seasonally adjust (X-11)	-2.05	0.12	1,789,041
4. Control for vacancy length	-2.73	0.19	447,387
5. Drop consecutive quarters	-1.58	0.12	772,619
6. Drop bonuses	-1.68	0.11	1,701,700
7. Occupation weights	-1.70	0.13	1,717,748
8. Region weights	-1.63	0.11	1,789,042
9. Time varying weights	-1.26	0.09	1,789,042
10. Pool across pay types	-1.68	0.11	1,788,591
11. Including ranges	-1.71	0.11	4,321,421

Notes: The first row reports the coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  from the baseline regression, that is, column 2 of Table 4. The second row reestimates the baseline regression at annual frequency. The third row seasonally adjusts unemployment and employment, at the state-quarter level for 1980–2021 data. The fourth row controls for the length of time for which vacancies are posted, taking the average by job and quarter. The fifth row excludes vacancies posted in the quarter immediately after another vacancy of the same job. The sixth row excludes vacancies with bonus pay. The seventh row reweights to target mean employment in each 6-digit occupation over 2010–2020. The eighth row reweights to target mean employment in each state over 2010–2021; the ninth row reweights to target employment in the previous quarter. The tenth row takes the mean wage across job titles, averaging over pay frequencies. The eleventh row includes vacancies posting a wage range, taking the mean of the range.

In all cases, the coefficient changes little or becomes more negative. Row 1 is our baseline specification from column 1 of Table 4. Rows 2 and 3 explore forms of seasonal adjustment, by reestimating at annual frequency in row 2; or seasonally adjusting unemployment and employment with the Census Bureau’s X-11 algorithm, in row 3. Rows 4 and 5 explore whether there are dynamic selection concerns, because jobs that repost vacancies without filling them may set wages differently from the others. Row 4 controls for the length of time for which vacancies were posted; row 5 excludes jobs that post vacancies in consecutive quarters, these jobs are most likely to repost without filling.<sup>17</sup> Row 6 verifies our results are robust to excluding jobs with bonus pay. Rows 7–9 study representativeness. Row 7 reweights to the occupation distribution at the 6-digit SOC level. Row 8 reweights to the regional distribution of employment with time invariant weights; row 9 does the same using time varying weights lagged by one quarter. Row 10 studies a broader definition of a job, as a job title within an establishment, while pooling across pay frequencies. Row 11 shows similar results if we also include wage ranges, and take the mean of the wage range.

Our Supplemental Appendix contains further robustness exercises. These exercises use different sources of variation to show that downward wage rigidity is pervasive. First, we find that downward wage rigidity is common across occupations and industries. Supplemental Appendix Table 7 estimates the baseline regression (1) for five broad occupations, and finds across all of them that wages are rigid downward and flexible upward. Supplemental Appendix Table 8 estimates the baseline regression for 10 broad industry sectors, and detects similar patterns across all sectors.

<sup>17</sup> Additionally, Supplemental Appendix Table 6 estimates the baseline regression separately for each quartile of the distribution of vacancy posting times, and finds similar results for each quartile.

We also find downward rigidity at the establishment level. In principle, even though job-level wages seem to be downwardly rigid, establishments might avoid wage rigidity by changing their mix of jobs. Therefore we reestimate the baseline regression (1) at the establishment level, by replacing the outcome variable with the mean nominal establishment wage, averaging across all jobs posted by an establishment in a given quarter. Supplemental Appendix Table 9 reports the results, finding a similar degree of rigidity.

One concern is that jobs posting only once—which do not enter the main sample of job-level differences—have relatively flexible wages. One test for this concern is to ask whether jobs that post infrequently, but more than once, tend to have more flexible wages. Supplemental Appendix Table 10 estimates downward wage rigidity, for jobs that post few or many times. Both kinds of jobs display similar of downward wage rigidity, suggesting that posting infrequently is not associated with more flexible wages.

Supplemental Appendix Table 11 estimates downward wage rigidity separately for each source of the vacancy data, since the data is drawn from a mix of online job boards and company websites. All sources show downward wage rigidity, suggesting our result is not driven by the mechanics of a particular method of posting vacancies. Supplemental Appendix Table 12 regresses wage growth, at the job level, on quarterly industry employment growth, at the 3-digit industry level. The regression shows that wages are rigid downward and flexible upward. Supplemental Appendix Figure 5 graphs the data corresponding to this regression, in a binned scatterplot, and finds patterns consistent with the regression. Supplemental Appendix Table 13 estimates our baseline regression by quartiles of establishment size, and finds that downward wage rigidity is pervasive across establishment size. Supplemental Appendix Table 13 also estimates our baseline regression across the four quartiles of the wage distribution, and again finds that downward wage rigidity is widespread.<sup>18</sup>

Now, we consider two qualifications to our finding of downward wage rigidity. Both are important, but we believe neither diminishes our results entirely. First, we investigate whether our results are explained by selection into posting wages. The main sample contains only jobs posting wage information. However, firms might post jobs without wage information—dropping from the regression sample—when they expect to renegotiate the wage during hiring. If true, then wage flexibility estimated from wage postings will understate true wage flexibility.

Therefore we estimate whether posting wages on vacancies is cyclical. To do so, we must study a different dataset from the main regression sample, to also include vacancies that do not post wages. For every job and quarter, we calculate the share of vacancies that post wages. We difference the share of vacancies posting wages at the job level, and then regress the job-level difference on unemployment changes.

<sup>18</sup>Supplemental Appendix Section D presents results that control for the length of time since the last posting. These controls alter the finding that wages are flexible upward; instead, wages respond little to decreases in unemployment. Conditioning on the length of time between vacancies, scatterplots in panels A–F of the Figure 10 suggest that wages do not respond to decreases in unemployment. The finding that wages are downwardly rigid is unaffected by these controls. These results are important for the reader to digest. However, as we discuss these covariates are a “bad control,” and should not be included in the regression (Angrist and Pischke 2009). The baseline results are robust using related control strategies that are not affected by the bad control problem.

Other than the outcome variable, the regression is similar to the baseline job-level regression (1).

Supplemental Appendix Table 14 reports the result and does not detect selection into posting wages. The main coefficients of interest are in column (2), which estimates the regression specification with time fixed effects. The coefficient on  $\Delta U_{st}$  is a precise zero. Therefore after a contraction, the share of vacancies posting wages does not change. The coefficient on  $\Delta U_{st} \times I(\Delta U_{st} < 0)$  is also zero, meaning that the share of vacancies posting wages does not change during expansions—whereas, if there were selection, the share of vacancies posting wages would rise during expansions.

Section V will take a more structural approach to the same issue, modeling the decision to post vacancies with wages and making a selection correction. This approach also does not suggest selection bias. In principle, there could be selection into posting wages at the state level, instead of the job level. However, in Supplemental Appendix Section C.2 we show that at the state level, the share of vacancies posting wages is acyclical. Therefore selection into wage posting at the state level also seems unlikely.

A second qualification is that our data measures the posted wage. Suppose that workers bargain instead of accepting the wage posted on the vacancy, so there is a gap between posted and realized wages. If this gap is cyclical, then posted wages may be less cyclical than realized wages, and our estimates will be biased. We suspect this concern—though reasonable—does not fully undermine our result.

To this end, we now show that posted wages are rigid even in occupations where wage bargaining is rare, or in occupations where workers typically accept the posted wage. We use data from the survey run by Hall and Krueger (2012), which asks workers first whether they bargained over wages, and second whether they accepted the posted wage. The survey also reports the 2-digit occupation of the new hire. We divide occupations into four quartiles, depending on whether either wage bargaining or accepting a posted wage is common. We estimate the baseline wage rigidity regression for every quartile.

This exercise is useful because of the following scenario: consider an occupation in which no worker bargains, or an occupation in which all workers accept the posted wage. Suppose we identify downward rigidity in posted wages for this occupation. There must also be downward rigidity in realized wages, because workers have accepted the posted wage and have not bargained. Our robustness test generalizes this idea by asking whether posted wages are still downwardly rigid in occupations where bargaining is unlikely, and accepting the posted wage is common.

Supplemental Appendix Table 15 reports the results. In panel A, we estimate wage rigidity separately for occupations in which wage bargaining is least common (column 1), through to occupations in which wage bargaining is most common (column 4). The coefficients are similar and do not seem to vary systematically across the groups. In panel B, we estimate wage rigidity separately for occupations in which accepting a posted wage is least common (column 1), through to occupations in which wage posting is most common (column 4). Again, the coefficients are similar and do not seem to vary systematically across the groups.

The results suggest that a cyclical gap between realized and posted wages cannot easily explain our findings. Posted wages are downwardly rigid even in occupations

where wage bargaining is rare or workers typically accept the posted wage. Moreover, wage rigidity is similar across all occupations, regardless of the degree of bargaining or accepting a posted wage.

#### IV. Job Composition and the Average Wage for New Hires

So far, we have provided new evidence that a measure of the wage for new hires is rigid downward and flexible upward. Previous work does not detect this asymmetry. However, previous work often studies the *average* wage for new hires, from worker-level survey data that controls for worker characteristics and averages across jobs. This section shows that due to job composition, average wages may have higher variance than job-level wages, meaning regressions with average wages have limited power to detect downward wage rigidity. We also discuss a related issue, which turns out to be less important during our sample period: job composition might create a form of omitted variable bias.

To start, we define measures of the wage for new hires at the average and job level. The wage for a newly hired worker in job type  $i$ , state  $s$ , and quarter  $t$  is  $w_{ist}$ . The share of new hires in job type  $i$  during the state-quarter is  $\nu_{ist}$ . Our dataset of wage postings measures growth in the *job-level* wage for new hires,  $\Delta \log w_{ist}$ . Some previous work, surveyed by Pissarides (2009), measures the *average* wage for newly hired workers, from survey data without information on jobs or establishments. Researchers study the change in the mean log wage  $\Delta \overline{\log w_{st}} = \Delta \left[ \sum_i \nu_{ist} \log w_{ist} \right]$ .<sup>19</sup>

Estimates using average wages often have large standard errors. To understand this pattern, first note that average and job-level wage growth can differ if job composition changes. A first order expansion of average wages yields an additive relationship between average wage growth, job-level wage growth, and composition:

$$(2) \quad \Delta \overline{\log w_{st}} \approx \sum_i \nu_{ist} \Delta \log w_{ist} + \sum_i \log w_{ist} \Delta \nu_{ist}.$$

Average wages can change, even if job-level wages do not change. Suppose that wages are unchanged at the job level during a given quarter—that is, the first term on the right hand side of equation (2) is zero. If the share of low-wage hires increases, wages change due to composition. The second term on the right hand side of equation (2) falls, so average wages fall.

From inspecting equation (2), we can see how job composition can raise the variance of average wages relative to job-level wages. Given downward rigidity, job-level wage growth  $\Delta \log w_{ist}$  is small. So the first term on the right hand side of equation (2) has low variance. However, suppose that the share of job types  $\nu_{ist}$  is volatile. Then the second term on the right hand side can have a high variance, so that average wage growth can have higher variance than job-level wage growth. Consider the example of an economy with high-wage bankers and low-wage baristas. Suppose that sometimes there are relatively many banker jobs, and sometimes

<sup>19</sup>Previous work often conditions on worker characteristics, by residualizing log wages against worker characteristics and then taking the average (e.g., Haefke, Sontag, and Van Rens 2013); or equivalently by running regressions on worker-level data while controlling for worker characteristics (e.g., Gertler, Huckfeldt, and Trigari 2020).

many barista jobs. Then average wages will vary, even if wages change little for either bankers or baristas.

As well as raising the variance of average wages, job composition might also create a form of omitted variable bias. Suppose that the share of low-wage jobs rises during recessions. Then, from equation (2), average wages must systematically fall during recessions, which confounds estimates of wage rigidity (Solon et al. 1994). However, in Supplemental Appendix Table 16, we find that the share of vacancy creation in low-wage jobs was not cyclical during this period, different from earlier periods (Hagedorn and Manovskii 2013; Gertler, Huckfeldt, and Trigari 2020). This finding is consistent with other work over the COVID-19 recession, which finds that declines in vacancy posting were widespread and similarly large across most occupations and industries (Forsythe et al. 2020).<sup>20</sup> Therefore we do not focus on omitted variable bias due to job composition, though it is likely important at other times.

We now explain how job composition affects inference, more formally. Our benchmark regression estimates downward rigidity using job-level wage variation. That is, we study the population regression function,

$$(3) \quad \Delta \log w_{ist} = \alpha + \gamma_t + \beta_{\text{Job Level}} \Delta U_{st} + \delta_{\text{Job Level}} I[\Delta U_{st} < 0] \Delta U_{st} + \varepsilon_{ist},$$

where  $\varepsilon_{ist}$  has bounded variance. We are interested in  $V[\hat{\delta}_{\text{Job Level}}]$ , the variance of the OLS estimator of  $\delta_{\text{Job Level}}$ . If  $\delta_{\text{Job Level}}$  is negative, then the wage for new hires is more rigid downward than upward at the job level.

Suppose a researcher only has access to average wages for new hires, as in prior work. A natural regression to study downward wage rigidity in average wages is

$$(4) \quad \Delta \overline{\log w_{st}} = \bar{\alpha} + \bar{\gamma}_t + \beta_{\text{Average}} \Delta U_{st} + \delta_{\text{Average}} I[\Delta U_{st} < 0] \Delta U_{st} + \bar{\varepsilon}_{st}.$$

This regression is the analogue of our job-level regression with average wages as the outcome variable. If estimates of  $\delta_{\text{Average}}$  are negative, then one concludes that average wages are downwardly rigid. If average wages are noisy, then the variance of the OLS estimator of  $\delta_{\text{Average}}$ , which we term  $V[\hat{\delta}_{\text{Average}}]$ , will be large.

In Supplemental Appendix Section E, we show that job composition inflates the variance of  $\hat{\delta}_{\text{Average}}$  relative to  $\hat{\delta}_{\text{Job Level}}$ . Thus regressions using average wages may lack the power to detect downward rigidity, even if it is present at the job level. From equation (2), the difference between job-level and average wage changes comes from changing job composition. In a regression with average wages, the residual variance is higher, creating noisier estimates. By contrast, regressions with job-level data purge noise due to job composition, and become precise.

We now show that in the data, job composition makes estimates with average wages imprecise. To understand this point, we compare estimates of  $\hat{\delta}_{\text{Average}}$  in equation (4) with our baseline estimate of  $\hat{\delta}_{\text{Job Level}}$ . For the outcome variable, we construct average posted wage at the state level, from Burning Glass; and the average wage for

<sup>20</sup>In contrast to the pattern for vacancy creation, separations were much larger for low-wage workers during the COVID-19 recession (Cajner et al. 2020).




FIGURE 5. ESTIMATES OF DOWNWARD RIGIDITY IN JOB LEVEL AND AVERAGE WAGES

Notes: The top row reports the estimate of  $\hat{\delta}_{\text{Job Level}}$ , which estimates downward rigidity with job-level data on posted wages from Burning Glass. The next three rows report various estimates of  $\hat{\delta}_{\text{Average}}$ , which estimates downward rigidity with average wages for new hires from Burning Glass. The second through fourth rows use, as the average wage measure, state-quarter average wages from Burning Glass, state-quarter average wages from the CPS controlling for worker information, and national average wages from the CPS. See Supplemental Appendix Table 17 for details.

new hires from the Current Population Survey (CPS). For average wages in the CPS, we follow the state-of-the-art procedure in Haefke, Sonntag and Van Rens (2013), by measuring averages after residualizing wages against 3-digit occupation and industry fixed effects, and demographic controls; we also include job switchers in the average as in Fallick and Fleischman (2004).

We study quarter-by-state data for 2010–2020:II, as in our benchmark regression. We report the standard error of  $\hat{\delta}_{\text{Average}}$  to contrast with the standard error of  $\hat{\delta}_{\text{Job Level}}$ . In both cases, we cluster standard errors at the state level. This procedure consistently estimates the standard deviation of the estimators  $\hat{\delta}_{\text{Average}}$  and  $\hat{\delta}_{\text{Job Level}}$ , given that the regressor  $\Delta U_{st}$  varies at the state level.

Figure 5 reports the standard error of downward wage rigidity estimates from job-level and average wages. The difference in precision between the estimates using average and job-level wages is enormous, meaning job composition does, indeed, inflate the variance of estimators of downward rigidity. The top row of Figure 5 reports the standard error of our job-level estimate of downward wage rigidity,  $\hat{\delta}_{\text{Job Level}}$ . The second row reports the standard error of  $\hat{\delta}_{\text{Average}}$ , the estimate of downward rigidity from average wages, using average posted wages from Burning Glass. The third row reports the standard error of  $\hat{\delta}_{\text{Average}}$ , with average wages for new hires from the CPS. The fourth row estimates  $\hat{\delta}_{\text{Average}}$  using national wage growth for new hires and national unemployment changes for 1984–2007. The sample period and measure of wages is the same as Haefke, Sontag, and Van Rens (2013). In all the regressions that use average wages instead of job-level wages, the standard error is far higher. Therefore the variance due to job composition is large in practice, and precludes researchers from detecting downward rigidity in average

wages. Supplemental Appendix Table 17 reports the point estimates and standard errors from the regressions in Figure 5.<sup>21,22</sup>

### V. Selection into Vacancy Posting

So far, we have exploited information on jobs, contained in vacancy postings, to estimate wage rigidity. However, selection into vacancy posting may affect our estimates. To understand such effects, this section studies a sample selection model following Heckman (1979), in order to apply a selection correction.<sup>23</sup>

In the model, time is discrete and there is a finite set of jobs. In period  $t$  and region  $s$ , job  $i$  decides whether to post a vacancy with a wage. The decision depends on a processes for aggregate unemployment  $U_{st}$  and an independent process for idiosyncratic productivity  $\phi_{ist}$ . The decision also depends on the latent wage  $w_{ist}^*$ —the wage that jobs would pay, if they were to post a vacancy. When the job posts a vacancy, the actual and latent wage are equal, so  $w_{ist}^* = w_{ist}$ . Otherwise the latent wage is not observed.

We focus on a log-linear relationship between latent wages, unemployment and idiosyncratic productivity

$$(5) \quad \log w_{ist}^* = \alpha - \beta U_{st} + \gamma \phi_{ist},$$

with  $\gamma \geq 0$ . Wages depend on aggregates, summarized by unemployment—and on the particulars of the job, summarized by idiosyncratic productivity.  $\beta$  is the measure of wage rigidity to be estimated. This equation resembles a simplified version of our benchmark regression, but replaces the regression residual with a structural idiosyncratic productivity term.

Jobs choose to post a vacancy if the value is high enough, given previous employment at the job  $n_{is,t-1}$ . The value of a vacancy is  $V_{ist} = V(U_{st}, \phi_{ist}, n_{is,t-1})$ , which is decreasing in its first argument, increasing in its second, and decreasing in its third. Higher unemployment signals unprofitable aggregate conditions that lower the value of a vacancy, whereas higher productivity raises the value. Higher past employment in the job may lower the value, due to decreasing returns to scale. The dependence on  $U_{st}$  and  $\phi_{ist}$  account for the indirect effect on vacancy value via latent wages.

In the model, jobs post vacancies when the value is positive. Therefore the vacancy posting decision is

$$(6) \quad \xi_{ist} = 1 \{V(U_{st}, \phi_{ist}, n_{is,t-1}) \geq 0\},$$

where  $\xi_{ist} \in \{0, 1\}$  is an indicator for whether the job posts vacancies. This simple sample selection model captures a key feature of the data: Jobs enter and exit

<sup>21</sup> Supplemental Appendix Table 17 also considers a measure of average wages that controls for worker fixed effects, constructed by Basu and House (2016).

<sup>22</sup> The job-level estimates not more precise simply because Burning Glass contains many jobs. The degrees of freedom in the regression is the number of states, and not the number of jobs, and standard errors are clustered at the state level.

<sup>23</sup> Martins, Solon, and Thomas (2012) provide an alternative method by estimating wage cyclicality for new hires only for a sample of “entry jobs,” which hire at all phases of the business cycle.

the vacancy data depending on idiosyncratic and aggregate shocks as well as wage rigidity.

Now, we use the simple model to make two points: first, regressions with average wages are biased by selection into vacancy posting; second, job-level regressions should reduce this bias.

To show selection bias, we take expectations of the wage equation (5) conditioning on a vacancy having positive value and being posted, which implies

$$(7) \quad E[\log w_{ist} | U_{st}, V_{ist} \geq 0] = \alpha - \beta U_{st} + \gamma E[\phi_{ist} | V_{ist} \geq 0].$$

We have a regression equation with an omitted variable bias. Average wages amongst jobs posting vacancies at time  $t$  is the outcome variable, and state unemployment is the regressor. The regression residual has a nonzero mean because the value of a vacancy  $V_{ist}$  depends on unemployment: This term represents selection bias. Estimates of  $\beta$  from regressing average wages on unemployment are always biased towards zero.

Equation (7) shows that due to selection, wages may seem rigid even if they are truly flexible. When aggregate unemployment  $U_{st}$  is high, then only jobs with high idiosyncratic productivity  $\phi_{it}$  post vacancies. However, when aggregate unemployment is low, low productivity jobs also post vacancies, which lowers average wages all else equal.

According to the model, regressions that difference wages at the job level should reduce selection bias. Taking a difference of equation (7) within the same job over any periods  $t$  and  $t - j$  in which the job posts a vacancy implies

$$(8) \quad \begin{aligned} E[\Delta_{t,t-j} \log w_{ist} | U_{st}, U_{s,t-j}, V_{ist} \geq 0, V_{is,t-j} \geq 0] \\ = \alpha - \beta \Delta_{t,t-j} U_{st} + \gamma E[\Delta_{t,t-j} \phi_{ist} | V_{ist} \geq 0, V_{is,t-j} \geq 0], \end{aligned}$$

which is a regression equation differencing wages at the job level. This equation resembles the baseline job-level regression (1), with a selection bias term. If  $\phi_{ist}$  is constant, so that idiosyncratic productivity does not vary over time, then the final term vanishes and there is no selection bias. If most of the variation in  $\phi_{ist}$  is time invariant, then the degree of residual selection bias will be small. Importantly, the residual bias in the job-level regression does not have a definite sign, unlike the bias in the regression with average wages: Selection in the job-level regression is now a more complicated function of both  $U_{st}$  and  $U_{s,t-j}$ , via their effects on  $V_{ist}$  and  $V_{is,t-j}$ .

Intuitively, the job-level regression should reduce selection bias because this regression compares the same job over time, and removes variation from a changing selection of jobs. In the particular case of constant idiosyncratic productivity, the effect of selection on the job's wages is the same in both periods, meaning that a difference of wages removes the effect of selection entirely.

Therefore in practice, job-level regressions will reduce selection bias if the relevant form of productivity is mostly time invariant at quarterly frequency. The available evidence tentatively supports this view. For example, Lachowska et al. (2022) study quarterly data from Washington state and find that 93 percent of the firm component of wages is explained by time invariant firm characteristics. However, this

evidence might not apply to our setting. We therefore turn to the data and a selection correction to see whether job-level regressions reduce selection bias.

### A. Selection Correction for the Job Level Regression

Now, we implement a standard Heckman (1979) correction, in order to test whether our job-level regressions are affected by selection bias.

To motivate the correction, observe that the vacancy posting decision (6) is a standard selection equation, which captures two ways in which past employment affects the likelihood of posting a vacancy. First, if there are decreasing returns to scale in employment, then the right hand side of the inequality in equation (6) falls as  $n_{is,t-1}$  increases. When jobs employ more workers, the value to hiring an additional worker by posting a vacancy may be lower. Second, jobs with different values of past employment, or jobs on different employment trends, will tend to have different conditional distributions of idiosyncratic productivity  $\phi_{ist}$ . For instance, jobs with growing employment may be receiving positive idiosyncratic shocks—and therefore be more likely to post vacancies.

The selection equation (6) suggests a standard Heckman two step procedure for correcting selection bias. In the first step, we estimate the probability of posting a vacancy, separately for periods  $t$  and  $t - j$ , as a function of past employment in the job. In the second step, we add estimates of the probability of vacancy posting in both periods as controls in our benchmark job-level regression, equation (8). Following the standard Heckman logic, these controls absorb the selection bias term in equation (8).

We use the Heckman estimator, proposed by Das, Newey, and Vella (2003), which does not impose parametric assumptions about selection, and flexibly controls for the effect of aggregate variables on selection. Supplemental Appendix Section F contains more details.

The intuition underlying the Heckman estimator is straightforward. Jobs select into vacancy posting when they receive idiosyncratic shocks. The Heckman estimator “backs out” the idiosyncratic shocks (using the relationship between whether a job posts a vacancy and its past employment) in order to create a control for the effect of selection on wages.

We motivated this estimator as a correction for bias from selection into vacancy posting. However, this estimator also corrects bias from selection into choosing to post a wage on a vacancy: since our model can be reinterpreted as a framework for choosing whether to post wages on a vacancy, given aggregate and idiosyncratic conditions.

To implement the Heckman correction, we need to observe the employment of the job, even when the job does not post vacancies: Burning Glass cannot provide this information by definition. In order to obtain selection-relevant information on jobs, we create a dataset that merges Burning Glass to annual establishment-level employment from Dun and Bradstreet (D&B). Dun and Bradstreet is a business analytics company that collect information on the universe of establishments in the United States from 1990–2020 inclusive, for the purpose of credit scoring and marketing. D&B collects data on employment at the start of the year, industry classification and establishment age. We do a fuzzy merge between establishments in Burning

Glass and D&B, matching on firm name, city, county and state (we describe our merge algorithm in Supplemental Appendix Section C.5). For cost reasons, we purchased data from D&B only on the largest 30 percent of employers in Burning Glass. We achieve a high merge rate; 75 percent of our main sample matches to establishments in D&B.<sup>24</sup>

Dun and Bradstreet is known to measure employment poorly for very small establishments. Once these establishments are removed, D&B appears to match official employment sources at the industry and regional level (Haltiwanger, Jarmin, and Miranda 2013; Barnatchez, Crane, and Decker 2017). Our extract from D&B does not include very small establishments, suggesting our extract should measure establishment employment reasonably well. Supplemental Appendix Section C.5 compares our D&B extract to official data from the Business Dynamics Statistics (BDS), and find a reasonably close match.

### B. Selection Corrected Estimates of Job-Level Regression

We apply our selection correction to the merged dataset of D&B employment and Burning Glass vacancies. Table 8 reports the results, and shows that selection bias is moderate in the job-level regressions. Column 1 reports our baseline job-level regression with time fixed effects, without controlling for selection bias, but restricting to the merged subsample. The estimates are similar to the full regression sample: Wages are rigid downward and flexible upward, with regression coefficients on  $\Delta U_{st}$  and  $I(\Delta U_{st} < 0) \times \Delta U_{st}$  of  $-0.05$  and  $-1.78$ , respectively. Column 2 corrects for selection. The results are broadly similar. The estimate of downward wage rigidity—the coefficient on the  $\Delta U_{st}$  regressor—becomes slightly closer to zero, with a value of  $-0.005$ . The estimate of upward wage flexibility—the coefficient on the  $I(\Delta U_{st} < 0) \times \Delta U_{st}$  regressor—also becomes closer to zero. The estimate is  $-1.48$ , a 17 percent change versus the estimate without the selection correction. Column 3 reports the difference in estimates between the regression coefficients with and without selection corrections, along with the standard errors of the difference.<sup>25</sup>

This selection correction suggests in practice the baseline job-level regression is not greatly affected by bias from selection into vacancy posting. If anything, selection biases the job-level regression towards estimating too little wage rigidity. However, our selection correction is based on strong assumptions. Notably, we require that lagged establishment employment does not enter the wage equation (5) other than through its relationship with idiosyncratic productivity, which is necessary for identification (Das, Newey, and Vella 2003).

To strengthen our finding of moderate selection bias, we close this section with a reduced-form approach to testing for selection. Equation (8) shows that selection bias will be strong for types of jobs that occasionally post vacancies. However, for

<sup>24</sup>Dun and Bradstreet was used by, for instance, Crouzet and Mehrotra (2020).

<sup>25</sup>Supplemental Appendix Table 18 carries out several robustness tests. We aggregate the regressions to the establishment-level and to the annual-level data; we use a simple control for employment growth instead of the full selection correction; we include entering and exiting establishments; we trim the propensity score for entry to values between zero and one; and we include lagged wages in the selection equation.

TABLE 8—CORRECTING SELECTION BIAS—NOMINAL POSTED WAGES AND UNEMPLOYMENT

	Nominal wage growth at the job level, $\Delta \log w_{jst}$		
	Without selection correction	With selection correction	Difference due to selection bias
$\Delta U_{st}$	-0.054 (0.014)	-0.005 (0.023)	0.049 (0.024)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.782 (0.097)	-1.477 (0.075)	0.305 (0.041)
Time fixed effect	✓	✓	✓

*Notes:* Column 1 estimates the baseline specification from column 2 of Table 4, restricting to the subset of the main sample that merges to Dun and Bradstreet. Column 2 reestimates the same specification but applies the selection correction outlined in the main text. Column 3 tests the null hypothesis that there is no selection bias. Each entry in column 3 reports the difference between the corresponding coefficients in panel A and B, and the standard error of this difference in parentheses. Counts refer to the number of differenced observations.

jobs that constantly post vacancies, regardless of aggregate or idiosyncratic conditions, selection bias will be small.

In practice, we can look for selection bias by comparing wage rigidity across occupations with low and high turnover. If wage rigidity were very different across occupations with high turnover—which are unlikely to be subject to selection bias—compared with low turnover occupations, then the latter group may be subject to selection bias.

Supplemental Appendix Table 19 shows that wage rigidity is similar across occupations with low and high turnover. We construct a measure of turnover as the ratio of Burning Glass vacancies to occupation employment from the Occupational Employment Statistics. We then estimate our baseline regression, equation (8), with time-fixed effects, for four quartiles of the occupational turnover distribution. The estimates, both of downward wage rigidity and upward wage flexibility, are similar across the four quartiles. By itself, this test cannot fully rule out selection bias. However, the test does suggest that selection bias is not a key feature of the data.<sup>26</sup>

## VI. Conclusion

We use a new dataset to show downward rigidity in a measure of the wage for new hires at the job level. We present two pieces of evidence. First, the posted wage rarely changes between successive vacancies at the same job. When wages do change for a given job, they are three times more likely to rise than to fall. These findings suggest a downward constraint on the wage in newly created jobs. Second, at the job level, the posted wage rises during expansions but does not fall during contractions, meaning wages are rigid downward and flexible upward at the job level.

One important question that our paper does not answer is *why* the wage for new hires might be more rigid downward than upward at the job level. Some plausible mechanisms for downward wage rigidity apply to continuing workers and not for new hires. For example, firms might offer implicit contracts in the form of downwardly

<sup>26</sup>One limitation of this test is that the Occupational Employment Statistics measures employment whereas Burning Glass measures vacancies; however official sources do not report disaggregated vacancy data.

rigid wages to continuing workers, and not extend the same insurance to new hires (Beaudry and DiNardo 1991). We hope future work can investigate the mechanisms behind downward rigidity for new hires, such as internal equity between continuing workers and new hires (Bewley 2002).

## REFERENCES

- Acemoglu, Daron, Amy Finkelstein, and Matthew J. Notowidigdo.** 2013. "Income and Health Spending: Evidence from Oil Price Shocks." *Review of Economics and Statistics* 95 (4): 1079–95.
- Acharya, Sushant, Julien Bengui, Keshav Dogra, and Shu Lin Wee.** 2018. "Slow Recoveries and Unemployment Traps: Monetary Policy in a Time of Hysteresis." Unpublished.
- Angrist, Joshua D., and Brigham Frandsen.** 2022. "Machine Labor." *Journal of Labor Economics* 40 (S1): S97–S140.
- Angrist, Joshua D., and Alan B. Krueger.** 1995. "Split-Sample Instrumental Variables Estimates of the Return to Schooling." *Journal of Business & Economic Statistics* 13 (2): 225–35.
- Angrist, Joshua D., and Jörn-Steffen Pischke.** 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Baqae, David, and Emmanuel Farhi.** 2022. "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis." *American Economic Review* 112 (5): 1397–1436.
- Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk.** 2014. "Some Evidence on the Importance of Sticky Wages." *American Economic Journal: Macroeconomics* 6 (1): 70–101.
- Barnatchez, Keith, Leland D. Crane, and Ryan A. Decker.** 2017. "An Assessment of the National Establishment Time Series (Nets) Database." Unpublished.
- Barnichon, Regis, Davide Debortoli, and Christian Matthes.** 2022. "Understanding the Size of the Government Spending Multiplier: It's in the Sign." *Review of Economic Studies* 89 (1): 87–117.
- Barro, Robert J.** 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3 (3): 305–16.
- Basu, Susanto, and Christopher L. House.** 2016a. "Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 297–354. Elsevier.
- Basu, Susanto, and Christopher L. House.** 2016b. Aggregate Wages for New Hires (supplement to NBER Working Paper No. 22279). NBER Working Paper Supplementary Materials. <https://www.nber.org/papers/w22279> (accessed August 11, 2025).
- Beaudry, Paul, and John DiNardo.** 1991. "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data." *Journal of Political Economy* 99 (4): 665–88.
- Bewley, Truman F.** 2002. *Why Wages Don't Fall During a Recession*. Harvard University Press.
- Bils, Mark, Marianna Kudlyak, and Paulo Lins.** 2022. "The Quality-Adjusted Cyclical Price of Labor." Unpublished.
- Bils, Mark J.** 1985. "Real Wages over the Business Cycle: Evidence from Panel Data." *Journal of Political Economy* 93 (4): 666–89.
- Cacciatore, Matteo, and Federico Ravenna.** 2021. "Uncertainty, Wages and the Business Cycle." *Economic Journal* 131 (639): 2797–2823.
- Cajner, Tomaz, et al.** 2020. "The US Labor Market during the Beginning of the Pandemic Recession." *Brookings Papers on Economic Activity* 2: 3–33
- Card, David, and Dean Hyslop.** 1997. "Does Inflation "Grease the Wheels of the Labor Market"?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer, 71–122. University of Chicago Press.
- Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal.** 2012. "Real wages and the business cycle: Accounting for worker, firm, and job title heterogeneity." *American Economic Journal: Macroeconomics* 4 (2): 133–52.
- Carnevale, Anthony P., Tamara Jayasundera, and Dmitri Repnikov.** 2014. *Understanding Online Job Ads Data*. Georgetown University.
- Center for Economic and Policy Research (CEPR).** 2025. *CPS Outgoing Rotation Group (ORG) Uniform Extracts*. CEPR Data. <https://ceprdata.org/cps-uniform-data-extracts/cps-outgoing-rotation-group/cps-org-data/> (accessed August 11, 2025).
- Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and Opportunity Insights.** 2020. *Economic Tracker: COVID-19 Policies and Outcomes (State-level)*. <https://github.com/OpportunityInsights/EconomicTracker>. Files used: COVID - State - Daily.csv, Policy Milestones - State.csv (accessed August 11, 2025).

- Chetty, Raj, John N. Friedman, and Michael Stepner.** 2024. “The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data.” *Quarterly Journal of Economics* 139 (2): 829–89.
- Chodorow-Reich, Gabriel, Olivier Darmouni, Stephan Luck, and Matthew Plosser.** 2022. “Bank Liquidity Provision Across the Firm Size Distribution.” *Journal of Financial Economics* 144 (3): 908–32.
- Chodorow-Reich, Gabriel, and Johannes Wieland.** 2020. “Secular Labor Reallocation and Business Cycles.” *Journal of Political Economy* 128 (6): 2245–87.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt.** 2016. “Unemployment and Business Cycles.” *Econometrica* 84 (4): 1523–69.
- Crouzet, Nicolas, and Neil R. Mehrotra.** 2020. “Small and Large Firms over the Business Cycle.” *American Economic Review* 110 (11): 3549–3601.
- Daly, Mary C., and Bart Hobijn.** 2014. “Downward Nominal Wage Rigidities bend the Phillips Curve.” *Journal of Money, Credit and Banking* 46 (S2): 51–93.
- Das, Mitali, Whitney K. Newey, and Francis Vella.** 2003. “Nonparametric Estimation of Sample Selection Models.” *Review of Economic Studies* 70 (1): 33–58.
- Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger.** 2013. “The Establishment-Level Behavior of Vacancies and Hiring.” *Quarterly Journal of Economics* 128 (2): 581–622.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh.** 1998. *Job Creation and Destruction*. MIT Press
- Dingel, Jonathan I., and Brent Neiman.** 2020a. “How Many Jobs Can Be Done at Home?” *Journal of Public Economics* 189: 104235.
- Dingel, Jonathan I., and Brent Neiman.** 2020b. *Work-from-Home Measures: state\_workfromhome csv and occupations\_workathome.csv*. <https://github.com/jdingelDingelNeimanworkathome> (accessed August 11, 2025).
- Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson.** 2020. “A Plucking Model of Business Cycles.” Unpublished.
- Eliaz, Kfir, and Rani Spiegler.** 2014. “Reference Dependence and Labor Market Fluctuations.” *NBER Macroeconomics Annual* 28 (1): 159–200.
- Fallick, Bruce, and Charles A. Fleischman.** 2004. “Employer-to-Employer Flows in the US Labor Market: The Complete Picture of Gross Worker Flows.” Unpublished.
- Forsythe, E., L. B. Kahn, F. Lange, and D. Wiczer.** 2020. “Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims.” *Journal of Public Economics* 189: 104238.
- Fukui, Masao.** 2020. “A Theory of Wage Rigidity and Unemployment Fluctuations with On-the-Job Search.” Unpublished.
- Ganong, Peter, Pascal Noel, and Joseph Vavra.** 2020a. “US Unemployment Insurance Replacement Rates During the Pandemic.” *Journal of Public Economics* 191: 104273.
- Ganong, Peter, and Pascal Noel.** 2020b. *UI Replacement Rates: State-level replacements and wages*. [https://github.com/ganong-noel/ui\\_rep\\_rate](https://github.com/ganong-noel/ui_rep_rate) (accessed August 11, 2025).
- Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante.** 2018. “Aggregate Recruiting Intensity.” *American Economic Review* 108 (8): 2088–2127.
- Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari.** 2020. “Unemployment Fluctuations, Match Quality, and the Wage Cyclicalities of New Hires.” *Review of Economic Studies* 87 (4): 1876–1914.
- Gertler, Mark, and Antonella Trigari.** 2009. “Unemployment Fluctuations with Staggered Nash Wage Bargaining.” *Journal of Political Economy* 117 (1): 38–86.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz.** 2021. “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data.” *American Economic Review* 111 (2): 428–71.
- Guerrieri, Veronica, Guido Lorenzoni, Ludvig Straub, and Iván Werning.** 2021. “Monetary Policy in Times of Structural Reallocation.” Unpublished.
- Guerrieri, Veronica, Guido Lorenzoni, Ludvig Straub, and Iván Werning.** 2022. “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” *American Economic Review* 112 (5): 1437–74.
- Haefke, Christian, Marcus Sonntag, and Thijs Van Rens.** 2013a. “Wage Rigidity and Job Creation.” *Journal of Monetary Economics* 60 (8): 887–99.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens.** 2013b. *CPS Wage Data for New Hires (wage-data\_all2560\_www.csv; PAYEMS.dta)*. <https://thijsvanrens.com/wage/> (accessed August 11, 2025).
- Hagedorn, Marcus, and Iourii Manovskii.** 2013. “Job Selection and Wages over the Business Cycle.” *American Economic Review* 103 (2): 771–803.
- Hall, Robert E.** 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review* 95 (1): 50–65.

- Hall, Robert E., and Alan B. Krueger.** 2012. "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-The-Job Search." *American Economic Journal: Macroeconomics* 4 (4): 56–67.
- Hall, Robert E., and Alan B. Krueger.** 2020. *Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search (Survey Data)*. <https://www.openicpsr.org/openicpsr/project/114257/version/V1/view> (accessed August 11, 2025).
- Hall, Robert E., and Paul R. Milgrom.** 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review* 98 (4): 1653–74.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda.** 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95 (2): 347–61.
- Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson.** 2022. "The Slope of the Phillips Curve: Evidence from US States." Unpublished.
- Hazell, Jonathon, and Bledi Taska.** 2025. *Data and Code for "Downward Rigidity in the Wage for New Hires"*. American Economic Association; distributed by Inter-university Consortium for Political and Social Research. <https://doi.org/10.3886/E226222V1>.
- Heckman, James J.** 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153–61.
- Keynes, John Maynard.** 1936. *The General Theory of Employment, Interest, and Money*. Springer.
- Klenow, Peter J., and Oleksiy Kryvtsov.** 2008. "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent US Inflation?" *Quarterly Journal of Economics* 123 (3): 863–904.
- Kudlyak, Marianna.** 2014. "The Cyclical of the User Cost of Labor." *Journal of Monetary Economics* 68: 53–67.
- Kurmann, André, and Erika McEntarfer.** 2019. "Downward Nominal Wage Rigidity in the United States: New Evidence from Worker-Firm Linked Data." Drexel University School of Economics Working Paper Series 1.
- Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury.** 2022. "Do Firm Effects Drift? Evidence from Washington Administrative Data." *Journal of Econometrics* 233 (2): 375–95.
- Martins, Pedro S., Gary Solon, and Jonathan P. Thomas.** 2012. "Measuring What Employers do about Entry Wages over the Business Cycle: A New Approach." *American Economic Journal: Macroeconomics* 4 (4): 36–55.
- McLaughlin, Kenneth J.** 1994. "Rigid Wages?" *Journal of Monetary Economics* 34 (3): 383–414.
- Moen, Espen R.** 1997. "Competitive Search Equilibrium." *Journal of Political Economy* 105 (2): 385–411.
- Nakamura, Emi, and Jon Steinsson.** 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.
- National Bureau of Economic Research.** 2025. *CPS Basic Monthly Files (harmonized extracts)*. NBER Data. [https://www2.nber.org/data/cps\\_basic.html](https://www2.nber.org/data/cps_basic.html) (accessed August 11, 2025).
- Pissarides, Christopher A.** 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (5): 1339–69.
- Shimer, Robert.** 2004. "The Consequences of Rigid Wages in Search Models." *Journal of the European Economic Association* 2 (2-3): 469–79.
- Shimer, Robert.** 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker.** 1994. "Measuring the Cyclical of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics* 109 (1): 1–25.
- US Bureau of Labor Statistics.** 1990. *Quarterly Census of Employment and Wages (QCEW): CSVs by Area*, Quarterly (1990Q1–2021Q4; plus SIC files). QCEW Downloadable Files. <https://www.bls.gov/cew/downloadable-data-files.htm> (accessed August 11, 2025).
- US Bureau of Labor Statistics.** 2016. *Occupational Employment Statistics (OES), 2016 National Data*. OES/OES Data Tables. <https://www.bls.gov/oes/tables.htm>. (accessed August 11, 2025).
- US Bureau of Labor Statistics.** 2025a. *Current Employment Statistics (CES): All Series Text Files*. BLS Public Data Files. <https://download.bls.gov/pub/time.series/ce/>. Primary file: ce.data.0.All-CESSeries (accessed August 11, 2025).
- US Bureau of Labor Statistics.** 2025b. *Consumer Price Index (CPI): cu/ text files*. BLS Public Data Files. <https://download.bls.gov/pub/time.series/cu/> (accessed August 11, 2025).