Downward Rigidity in the Wage for New Hires

Jonathon Hazell Bledi Taska*

November 10, 2020

Downward wage rigidity is central to many explanations of unemployment fluctuations. In benchmark models, the wage for new hires is key, but there is limited evidence of downward rigidity on this margin. We introduce a dataset that tracks the wage for new hires at the *job level*—across successive vacancies posted by the same job title and establishment. We show that the wage for new hires is rigid downward but flexible upward, in two steps. First, the nominal wage rarely changes at the job level. When wages do change, they fall infrequently. Second, when unemployment rises, wages do not fall—but wages do rise strongly as unemployment falls. We show prior strategies cannot detect downward rigidity due to job composition. Then with a standard model, we argue downward wage rigidity at the job level is key for unemployment fluctuations. Unemployment responds four times more to negative than to positive labor demand shocks.

*First draft: July 2018. Hazell: Princeton University & London School of Economics (email: hazell@princeton.edu); Taska: Burning Glass Technologies. Hazell thanks Iván Werning, Daron Acemoglu, Jonathan Parker and Emi Nakamura for invaluable support and guidance. We also thank for their comments Martin Beraja, Sydnee Caldwell, Tyler Cowen, Juliette Fournier, John Grigsby, Erik Hurst, Gwen Jing, Mazi Kazemi, Eben Lazarus, Chen Lian, Christian Moser, Christina Patterson, Claudia Sahm, Pari Sastry, Anna Stansbury, Jón Steinsson, Emil Verner 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 and CREI. 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".

1 Introduction

Suppose there is downward wage rigidity—that is, 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). So, wage rigidity leads to large unemployment fluctuations over the business cycle (Hall, 2005a). Downward wage rigidity for *new hires* is particularly important (Pissarides, 2009). Employment is a long term contract. So, the present value of wages, which is tied to the wage for new hires, matters to workers and firms (Barro, 1977). Even if wages in continuing jobs change little, the present value of wages can still vary if the wage for new hires is flexible.

But consensus on wage rigidity for new hires is elusive. Job composition is a key challenge. Prior work often uses worker level survey data without job level information. This work studies the average wage for new hires, averaging over the jobs into which workers are hired, while controlling for worker characteristics. Pissarides (2009) surveys this work. If job composition varies over time, average wage changes reflect either changing job composition, or wage changes for individual jobs (Gertler and Trigari, 2009). As an example, consider an economy of high wage bankers and low wage baristas. Suppose the share of barista hires increases during recessions. Then average wages for new hires fall, even if wages fall for neither baristas nor bankers. Conversely, suppose the share of barista hires decreases during recessions. Then average wages will not fall, even if wages do fall for both barista and banker jobs.[1] Estimates in the prior literature are often too imprecise to draw conclusions. For example, the point estimate in Haefke, Sonntag, and Van Rens (2013) suggests strong procyclicality, but the confidence interval includes zero cyclicality.

There is evidence of downward wage rigidity for continuing workers (Grigsby, Hurst, and Yildirmaz, 2018). But plausible mechanisms predict wages are more flexible for new hires than for continuing workers. For example, firms might not cut wages to preserve workers' morale (Bewley, 2002). This consideration may matter less for new hires.

We study a dataset on the wage for new hires, from online vacancies, collected by Burning Glass Technologies. Our data has job and establishment level information on wages. Our main contribution is to show that the nominal wage for new hires is rigid downward, but flexible upward. We isolate job-level wage changes, to purge the effects of job composition. Then with a model we argue that wage rigidity for new hires, at the job level, is important for unemployment fluctuations.

[1] For simplicity we refer to previous work as studying the average wage for new hires. As we discuss later in more detail, previous work often studies the average wage for new hires *conditional on worker characteristics*, by controlling for worker level observables, worker fixed effects or proxies for job composition. Conditional on these controls, residual job composition may still affect average wages for new hires.

Our dataset contains wages on new vacancies, with job titles, establishment identifiers, and pay frequency, for 10% of all vacancies posted in the United States during 2010-2016. The dataset collects vacancies from online job boards and company websites. The dataset has limitations: it is not a representative sample, and records wages posted on vacancies instead of the realized wage paid to new hires. Still, the dataset seems to measure the wage for new hires. Average wages in Burning Glass closely track state-by-quarter measures of the average wage for new hires from both survey and administrative data.[2]

Our dataset has an advantage relative to prior work that measures the average wage for new hires. We can track *job level* variation in the wage for new hires—that is, the wage across successive vacancies posted by the same job title and establishment.[3] 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. Workers are typically hired once as a barista. So worker-level data cannot easily track the wage across successive workers hired as a barista at the Starbucks. By studying job level wages, we can purge wages changes due to job composition, which could obscure wage rigidity.

In the main contribution of the paper, we show that the nominal wage for new hires is rigid downward, but flexible upward. We have three findings. First, we detect signs of a constraint on wage setting for new hires. The wage for new hires rarely changes between successive vacancies at the same job, wages typically change once during 5 quarters. When wages do change for a given job, they are three times more likely to rise than to fall.

Second, at the job level, the wage for new hires rises during expansions but does not fall during contractions. Figure 1 illustrates the result. In the figure, the wage for new hires is averaged by job and quarter. On the y-axis is wage growth between two consecutive vacancies for the same job, from Burning Glass. On the x-axis is the growth in quarterly state level unemployment between the quarters in which the vacancies are posted, from the Bureau of Labor Statistics.[4] As state unemployment decreases, the wage for new hires rises strongly. As state unemployment increases, wages do not fall. Figure 1 isolates job-level wage growth for new hires. We remove variation from changing job composition, which might obscure downward wage rigidity. We study state unemployment to overcome the short time series dimension of the data. We then confirm the finding with regressions.[5]

[2] Posting vacancies on job boards or company websites is normally expensive, and vacancies are typically active for a month or less. We suspect these features discourage "stale" information in vacancies.

[3] Here, a "job" is a job title at an establishment.

[4] Since many jobs do not post in consecutive quarters, sometimes the fall in unemployment between postings is relatively large.

[5] Our results are robust to studying establishment variation instead of job titles within establishments; reweighting to the occupation and region distribution of jobs; studying annual frequency; using city, industry, or state-

Figure 1: Wage Growth for New Hires and Quarterly State Unemployment Changes

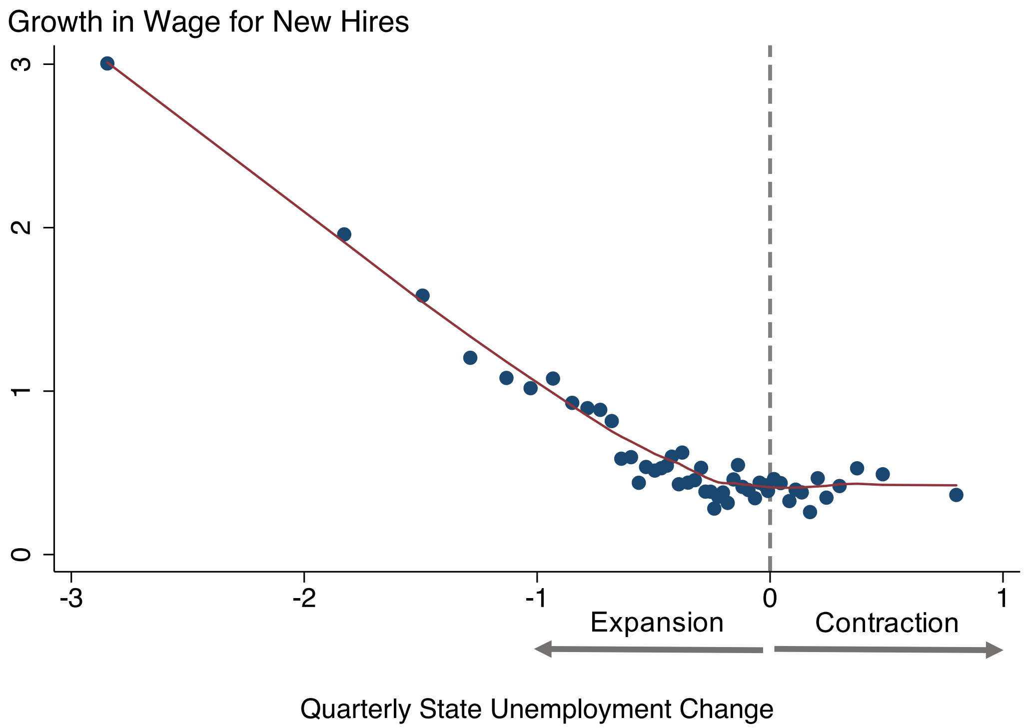

Notes: the graph plots binned wage growth for new hires, from Burning Glass, and binned state by quarter unemployment changes, from the Bureau of Labor Statistics. To construct wage growth, we take the mean wage within each job and quarter, and then take log differences at the job level. We use 50 bins, partial out time fixed effects, and add a non-parametric regression line.

Third, wage flexibility *upward* displays a form of state dependence consistent with downward rigidity. When there has been a contraction in the recent past, then subsequently, the wage for new hires responds little as unemployment falls. When there has been an expansion in the recent past, then subsequently, the wage for new hires responds strongly as regional unemployment falls. This state dependence is consistent with downward rigidity. Wages are "trapped too high" when there has been a contraction in the recent past, and a subsequent marginal increase in labor demand from the trough of the contraction does not raise wages. After an expansion in the recent past, wages overcome the downward constraint. This state dependence potentially explains the puzzle of "missing wage growth" during the early recovery from the Great Recession.[6]

These findings are new and perhaps surprising. To our knowledge, no previous work doc-

industry variation; for real wages; and after using an instrument for regional labor demand based on oil price shocks.

[6]See, for example, Federal Reserve Bank of Atlanta (2014). In 2014, the rapid decline in unemployment and slow growth of wages was deemed puzzling by many observers.

uments asymmetry or state dependence in the rigidity of new hire wages, nor infrequent wage changes for new hires at the job level. But previous work finds the same patterns for continuing wages (e.g. Kurmann and McEntarfer, 2017; Grigsby, Hurst, and Yildirmaz, 2018).[7] Our results suggest mechanisms that impose parity between the wage of new hires and continuing workers, such as internal equity (Bewley, 2002). Our paper supports an influential conjecture, from Gertler and Trigari (2009), that new hires' wages are just as rigid as continuing workers' wages.

Job level data is important for uncovering the results. The average wage for new hires, the object of previous studies, shows no sign of downward rigidity. We examine the average wage for new hires in our dataset. The average wage does not respond differently to rises versus falls in regional unemployment. Similarly, measures of the average wage from worker-level survey data, used in Haefke, Sonntag, and Van Rens (2013) and Basu and House (2016), do not display downward wage rigidity.

We find that job composition raises the variance of average wages. So, 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. So, average wages are also volatile. Then standard errors from regression with average wages are almost four times larger than counterparts using job-level wages. Previous work argues that composition biases point estimates (Gertler, Huckfeldt, and Trigari, 2016). We study a complementary effect: composition raises the variance of estimates.

We then develop a simple extension of the canonical labor search model. The model shows that measuring job-level wages is important because they govern unemployment fluctuations. In the model there are multiple types of job, in order to have a precise definition of the job-level wage. Otherwise, our model follows the standard Diamond-Mortensen-Pissarides framework.

We have two results that link our estimates of wage rigidity to unemployment fluctuations. First, we show analytically that changes in the wage for new hires, at the job level, are particularly important for unemployment fluctuations. So according to our model, we are measuring the right object in the data. Our model also reaffirms the canonical importance of the wage for new hires (Pissarides, 2009).[8] Second, we calibrate the model to gauge the quantitative implications of our estimates of wage rigidity. The degree of downward wage rigidity in the data means

[7] Grigsby, Hurst, and Yildirmaz (2018) study both new hire and continuing wages, we relate their findings on new hire wages to ours below. See Card and Hyslop (1997), Barattieri, Basu, and Gottschalk (2014), Daly and Hobijn (2014), and Makridis and Gittleman (2019) for further evidence of continuing wage rigidity. Continuing workers' wages do sometimes fall (Elsby and Solon, 2018; Jardim, Solon, and Vigdor, 2019). Similarly, we find that wages for new hires do sometimes fall at the job level.

[8] Continuing wages do matter in other labor search models, such as models with financial frictions (Schoefer, 2015), endogenous separations (Mortensen and Pissarides, 1994; Elsby and Michaels, 2013), on the job search (Menzio and Shi, 2011) or variable effort (Bils, Chang, and Kim, 2014).

that unemployment responds almost four times as much to negative labor demand shocks as to positive shocks. Moreover our model predicts large unemployment fluctuations on average, similar to US data. So, the wage rigidity in the data solves the "unemployment volatility puzzle" of Shimer (2005).

1.1 Related Literature

This paper contributes to three literatures. First, we contribute to the literature investigating the causes 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. Hall (2005a), Hagedorn and Manovskii (2008), Hall and Milgrom (2008) and Gertler and Trigari (2009) show that adding wage rigidity to the model leads to unemployment fluctuations as large as in US data. Pissarides (2009) emphasizes that in this model, the relevant wage is for newly hired workers. Our contribution to this literature is to show that the wage rigidity in the data generates large unemployment fluctuations.

This paper contributes to a second literature that measures wage rigidity for new hires. In the seminal paper, Bils (1985) regresses the wage for new hires on unemployment to measure wage cyclicality. The wage for new hires is from survey data, 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. Point estimates suggest strongly procyclicality, but confidence intervals often include weak or zero procyclicality. Gertler and Trigari (2009) emphasize the challenge of job composition in interpreting these results.[9]

Our work complements two papers that study wage rigidity for new hires and correct for job composition.[10] First, Gertler, Huckfeldt, and Trigari (2016) study wages for workers newly hired from unemployment. The average wage of workers hired from unemployment is plausibly less affected by job composition than the average wage of workers switching jobs. Gertler et al find weakly procyclical wages for workers hired from unemployment. Second, contemporaneous with the first draft of this paper, Grigsby, Hurst, and Yildirmaz (2018) use high quality payroll data on workers switching jobs, to measure the wage for new hires. They control for the effect of job composition on job switchers' wages using a matching estimator. This estimator matches workers switching jobs to similar workers who are not switching jobs. With this adjustment, the

[9] Solon, Barsky, and Parker (1994) and Grigsby (2019) emphasize the difficulty in interpreting wage cyclicality when worker composition can change.

[10] Other important papers using US data include Haefke, Sonntag, and Van Rens (2013), Hagedorn and Manovskii (2013), Kudlyak (2014), Basu and House (2016) and Doniger (2019).

wage for new hires is weakly procyclical.[11]

Our paper complements these two papers in two respects. First, our data is at the job level instead of the worker level. So, we can directly correct for job composition. The worker level data in the two related papers may not fully eliminate the effects of job composition. Second, we find that wages are rigid downward, but flexible upward; display state dependence; and change infrequently at the job level. The two related papers do not detect these patterns in the wage for new hires—though Grigsby, Hurst, and Yildirmaz (2018) do detect such patterns in continuing workers' wages.

Our paper contributes to a third literature, that studies the consequence of downward rigidity for asymmetries in unemployment. For example, Dupraz, Nakamura, and Steinsson (2016) shows that if wages for new hires are rigid downward and flexible upward, then unemployment rises sharply during contractions and falls more slowly during expansions.[12] Our paper provides the first evidence in the US that wages for new hires are rigid downward but flexible upward.

2 Data

We study an establishment level dataset of wages for new vacancies, with job titles, covering 2010-2016. The dataset was developed by Burning Glass Technologies, and draws from company websites and online job boards. The vacancy data contains wages and occupation information at the 2- 4- or 6-digit SOC code level.[13]

The dataset covers approximately 10% of vacancies posted in the US, either online or offline (Carnevale, Jayasundera, and Repnikov, 2014). Burning Glass draws from the near-universe of job vacancy postings, from 40,000 distinct online sources. No more than 5% come from any one source. The company employs a deduplication algorithm, to avoid double counting vacancies that post on multiple job boards.

The dataset contains detailed information on the wage in new vacancies. 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 attached to the vacancy. The hours measure is an important advantage. In the United States, administrative data typically does not contain hours worked, though it is

[11] Outside the US, Martins, Solon, and Thomas (2012), Carneiro, Guimarães, and Portugal (2012), Kaur (2019), Schaefer and Singleton (2019) and Choi, Figueroa, and Villena-Roldán (2020) study wage rigidity for new hires.

[12] Petrosky-Nadeau and Zhang (2013), Chodorow-Reich and Wieland (2017), Petrosky-Nadeau, Zhang, and Kuehn (2018), Acharya, Bengui, Dogra, and Wee (2018) and Cacciatore and Ravenna (2020) also study asymmetries in unemployment dynamics, when wages for newly hired workers are rigid downward and flexible upward.

[13] A 6 digit SOC code is granular—at the detail of, for example, a high school Spanish teacher.

available for some smaller states such as Washington and Minnesota. Survey data tend to have measurement error in wages (Bound and Krueger, 1991). The dataset also records any education requirements associated with the vacancy, such as high school diploma or undergraduate degree, if they are present.

The data report establishment and job title. 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. Throughout the paper, we use the term "job" to refer to a job-title within an establishment whose wages are paid at a given frequency (e.g. annual or daily).

The dataset overweights certain occupations that disproportionately post online. Appendix Figure 2 plots the relative share of Burning Glass occupations versus the 2014-2016 Occupational Employment Statistics (OES). In robustness exercises for our empirics, we reweight to the occupational or regional distribution of jobs, and find little change to our results.

Importantly, Hershbein and Kahn (2016) show that the representativeness of Burning Glass is stable over time at the occupation level. Though Burning Glass under-represents some occupations relative to the CPS, the *degree* to which these occupations are under-represented does not change. Hershbein & Kahn construct the share of new jobs in each 3 digit occupation, in both Burning Glass and the CPS. The occupations that are underweight in Burning Glass at the start of the sample period, are typically underweight by the same amount at the end of the sample period. Hershbein and Kahn's Online Appendix Figure A3 reports this result. By contrast, the accuracy of other popular online vacancy data, such as the Help Wanted Online series, is declining (Cajner and Ratner, 2016).

Table 1 reports summary statistics. There are many vacancies within each state-quarter. The dataset covers almost all 6-digit SOC occupations. A large fraction of jobs contain establishment and job title identifiers. Roughly half of the vacancies with wage information post a range of salaries. The rest post a point salary. For jobs that post a range, we use the mean of the range. Appendix Section C.6 explores in detail alternative ways of treating jobs that post a range, and finds that they do not make a difference to our key results.

The dataset of wages is a subset of the online vacancies provided by Burning Glass. In total, Burning Glass covers around 70% of vacancies in the United States (Carnevale, Jayasundera, and Repnikov, 2014). But only 17% of vacancies in Burning Glass include wages. So, the subset of vacancies in Burning Glass that include wages, which is the main sample that we study in the paper, is roughly 10% of total US vacancies. It is not clear why a minority of firms include wages on their vacancies. Marinescu and Wolthoff (2016) show that the decision to include wages in vacancies is a time-invariant characteristic of certain types of firms. These considerations are likely not relevant for business cycles: in Appendix Table 1, we show that firms' decisions to include wages in vacancies are not cyclical.

In many specifications, we study regional business cycle variation. We use quarterly unemployment from the Local Area Unemployment Statistics (LAUS) and state employment from the Quarterly Census of Employment and Wages (QCEW).

2.1 Measuring the Wage for New Hires with Burning Glass

We show that Burning Glass seems to measure the wage for new hires at business cycle frequency. We compare Burning Glass wages to the best available survey and administrative data on the wage for new hires and find that Burning Glass wages track these other measures. This step is important because of two key limitations of the data. First, Burning Glass records wages posted on vacancies, and not realized wages paid at the start of the hire. Second, Burning Glass is neither a representative sample, nor a census, of the wage for new hires. This finding sets the stage for our main empirical results: we can use the special features of our dataset to investigate wage rigidity for new hires. Henceforth we refer to Burning Glass wages simply as the wage for new hires.

First, we construct an alternative measure of the wage for new hires from the Current Population Survey (CPS), at the state-by-quarter level for 2010-2016. The wage for new hires is from workers switching jobs over the previous quarter, or entering jobs from unemployment. We use the rotating panel component of the CPS's basic monthly files, and wage data from the CPS Outgoing Rotation Group, following Haefke, Sonntag, and Van Rens (2013). Wages are usual hourly earnings for hourly and non-hourly workers.

We regress log CPS wages on log wages from Burning Glass, also at the state-quarter level. 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 halve the data in each state-quarter and calculate average state-quarter wages in each sub-sample. We then instrument for wages in one sub-sample 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.

Table 2, Panel A, reports the regressions, and Appendix Figure 1 presents a binned scatterplot of the regression. The regression coefficient is near one. So, when Burning Glass wages change by one percentage, wages for new hires from the CPS also change by roughly one percent. The results are similar if we add state or time fixed effects. Our estimates are fairly precise, and we cannot reject that the regression coefficient is 1. Thus the Burning Glass and CPS measures of the wage for new hires comove one-for-one—Burning Glass closely tracks other measures of the new hire wage. When restricted to the sample containing job identifiers, which form much of the analysis that follows, our estimates are virtually unchanged. Despite small sample sizes and measurement error in CPS data, the large Burning Glass dataset lets us obtain precise estimates.

We next compare Burning Glass wages to average earnings for newly hired workers, from administrative data at the state-quarter level for 2010 to 2016. This measure is administrative, from the Quarterly Workforce Indicators (QWI), and does not suffer from the small samples or measurement error in reported wages. However, the data reports earnings for new hires—inclusive of both hours worked and hourly wages—and cannot isolate a measure of hourly wages. We regress log state-quarter earnings for new hires, in the QWI, on log wages from Burning Glass, also at the state-quarter level. As before, we split the Burning Glass sample, and instrument for one half of the sample with the other.

Table 2, Panel B, reports the regressions. The elasticity of new hire earnings with respect to Burning Glass wages is near, but above one. After one percent of growth in Burning Glass wages, QWI earnings for new hires grow by 1.25 percent. The larger movement in QWI earnings than in Burning Glass wages likely reflects a positive comovement between hours and wages in the QWI, so that QWI earnings increase by more than wages. The results are similar after adding state or time fixed effects, and the estimates are again fairly precise. Again, when restricted to the sample containing job identifiers, which form much of the analysis that follows, our estimates are virtually unchanged. So, reassuringly, two different measures of the wage for new hires with different shortcomings and advantages, match the Burning Glass measure of wages. Appendix Table 2 compares Burning Glass wages to occupational and regional wages, and again finds a close match.

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% and as many as 80% 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.[14] 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, by sales, was $419 in 2017.[15] Companies posting on their own websites typically pay monthly fees to subcontractors. 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. Online job boards typically remove vacancies after one month, or request a

[14] Hall and Krueger find that 30% of workers knew their exact wage before being hired, and above 80% of workers "knew exactly or had a pretty good idea" of their wage before being hired. Meanwhile 35% of workers bargain over wages.

[15] See https://blog.proven.com/how-much-to-post-a-job.

Figure 2: An Example of a Job

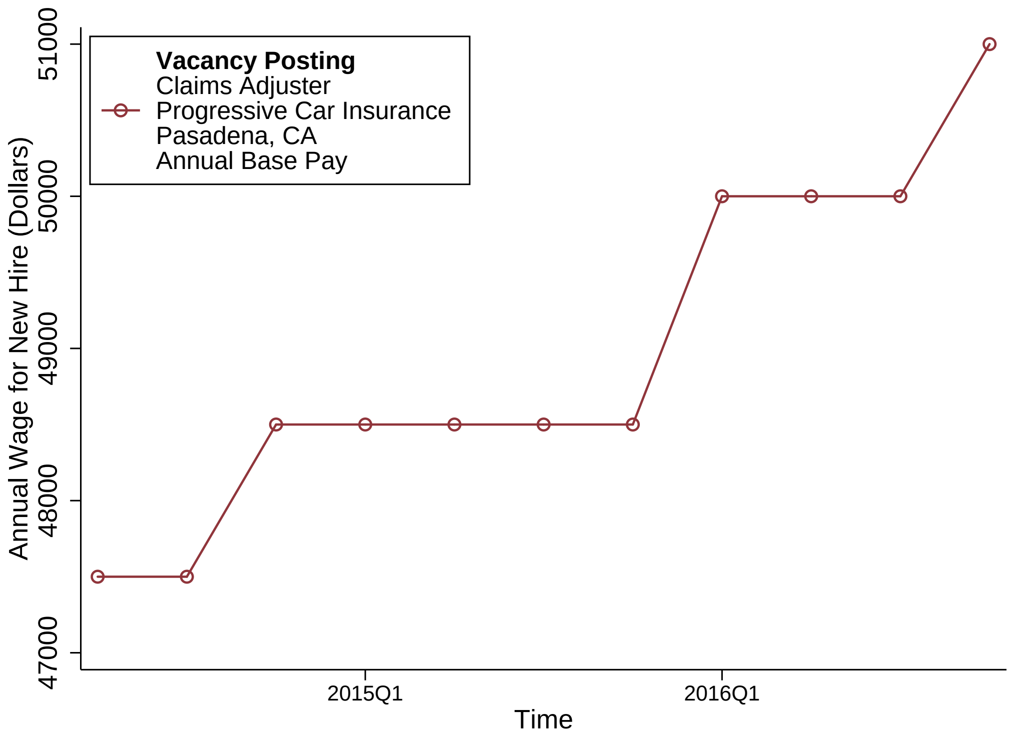

Notes: A job is a job-title by establishment by salary type by pay frequency unit, from Burning Glass. Claims Adjuster is a job title, for a vacancy posted by an establishment of Progressive Car Insurance, in Pasadena, California, for an annual base pay salary.

further fee for the vacancy to remain open. On company websites, the median duration of vacancies is 21 days, and 92% of vacancies are removed within the quarter.[16]

So, data from survey and administrative data therefore confirm that Burning Glass wages are a reasonable measure of the wage for new hires. We now explain what differentiates us from prior datasets.

2.2 Job-Level Data on the Wage for New Hires

Our dataset has a particular advantage not shared by prior data that measures the average wage for new hires. We can track the wage for new hires at 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 the wage for new hires faced by establishments.

Figure 2 displays job-level variation. We present a job that posts multiple vacancies. The

[16]The duration of vacancies is similar to the mean vacancy duration reported in Davis, Faberman, and Haltiwanger (2013) from the BLS's JOLTS survey, of 20 days.

firm is Progressive Car Insurance. The establishment is the branch of the firm in Pasadena, California. The job title is claims adjuster. The salary is an annual wage, base pay. When the vacancy posts multiple times within the quarter, we take the average. Then according to our definition, a job is a claims adjuster at the Pasadena establishment of Progressive Car Insurance. The job posts 11 vacancies over three years. We can track the wage across these vacancies—that is, we can track job-level changes in the wage for new hires. We can also track establishment level wage changes. We can study how wages change for the establishment of Progressive Car Insurance, pooling across all the jobs into which they hire workers in a quarter.

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. But we will argue that job level data has two further and related benefits. First, we will show that job level data can control for wage changes due to job composition. Without these controls, one cannot arrive at our new findings. Second, we argue that in standard models, variation in the wage for new hires at the job level is key for unemployment fluctuations. So, we are measuring a particularly important object in the data.

3 Constraints on Wage Setting for New Hires

This section presents evidence of a constraint on wage setting for new hires. The wage for new hires rarely changes between successive vacancies at the same job. When wages do change for a given job, they rarely fall. The job in Figure 2 also shows these patterns: the wage changes infrequently across vacancies, with three changes and no decreases over eleven vacancies and three years.

3.1 Hazard Estimation of the Probability of Wage Changes

We start by studying how often wages change, rise and fall at the job level.

First, we explain our treatment of the data. We aim to study wages across successive vacancies for the same job, and so restrict to jobs that post multiple vacancies. We take the mean wage for vacancies within each job-quarter. After these steps, there are roughly 1.6 million observations. Table 3 presents summary statistics for this subsample. There remains a large number of jobs for which we observe repeat postings. These jobs cover 99% of 6-digit SOC occupations in the US economy by employment share, and are well represented in all states. In robustness exercises, we will reweight at a fine level, to target the occupational or geographic distribution of jobs in the US, and find our main results unaffected.

We next confront a measurement challenge. We only observe wages for the quarters in which jobs post vacancies—wages are "missing" in other quarters. Therefore we cannot directly observe the probability that the wage for new hires changes, nor the duration of time for which wages are unchanged. We adapt a standard approach from the price setting literature to overcome this problem, first developed in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008).

We treated 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.[17] 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.[18] One can easily adapt this process to separately estimate the probability of wage increase and decrease. One can assume a constant hazard of wage increase or decrease, and estimate this process using the observed sequence of wage increases or decreases.

We use implied durations to measure for how long wages are unchanged, as in the price setting literature. Other simple procedures for calculating duration are biased downwards in the presence of left-censored spells (Heckman and Singer, 1984).

3.2 Infrequent Wage Changes at the Job Level

We find that the nominal wage for new hires changes infrequently, implying a constraint on wage setting at the job level.

[17]Appendix section C.1 presents tests of the constant hazard assumption and finds that it is approximately satisfied.

[18]We assume the hazard rate of the latent wage change is constant across time and common across all jobs within each 2 digit SOC occupation. Let $\{w_{it}\}$ be the sequence of log wages for job i and quarter t. Let γ_{it} be the gap in quarters between the wage at t and wage in the previous vacancy that was posted. Let I_{it} be an indicator for whether the wage changed, where $I_{it} = 1$ if $w_{it} \neq w_{i,t-\gamma_{it}}$. The quarterly hazard rate of wage change, assumed to be time-invariant, is given by λ, which we estimate by maximum likelihood. The likelihood function is $L = \prod_i \prod_t \left(1-e^{-\lambda\gamma_{it}}\right)^{I_{it}} \left(e^{-\lambda\gamma_{it}}\right)^{1-I_{it}}$. The probability of a wage change for each occupation is $f = 1 - e^{-\lambda}$. The implied duration of time for which a wage is unchanged is $d = 1/\lambda$. The overall probability of wage change is the median probability across occupations, weighted by the number of vacancies in each occupation. Similarly, the overall implied duration is the the weighted median of the implied duration for each occupation. We discard left-censored wage spells.

Table 4 reports the results. Across all columns, the probability of wage change is similar, and low—the corresponding implied durations are 5-6 quarters. Column (1) estimates the quarterly probability of wage change according to our method. Column (2) reweights vacancies at a granular level, to target the distribution of jobs from the 2014-6 Occupational Employment Statistics, the nationally representative establishment survey of occupational employment. Column (3) reweights to target the regional distribution of jobs from the QCEW. Column (4) drops jobs from the bottom quartile of the wage distribution, since minimum wages might cause infrequent changes. Results are similar in all cases, confirming that the wage for new hires changes infrequently at the job level. In Appendix Section B, we document the same statistics at annual frequency. The results are similar, again showing infrequent changes.

The wage for new hires changes infrequently even as many successive vacancies are posted for a given job. 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% of vacancies, less than 5 quarters have elapsed since the job posted a previous vacancy. So, 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 the wage for new hires already suggest a constraint on wage setting at the job level. We now show asymmetry—this constraint matters more for preventing wage falls than for wage rises.

3.3 Asymmetry in Wage Changes at the Job Level

At the job level, conditional on a wage changing, wages in new hires are more likely to rise than to fall. There is a downwards constraint on wage setting—while wages are more able to increase.

Figure 3 plots the distribution of wage growth, after removing vacancies with zero wage change. There are two clear points. First, conditional on a wage changing, wages in new hires rise more often than they fall. Secondly, wages "pile up" against the constraint—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 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. However, not all jobs post in consecutive quarters. We truncate the plot at $\pm 10\%$ wage growth.

We then estimate the probability of wage increases and decreases for new hires. The results

Figure 3: Distribution of Wage Growth for New Hires, Conditional on Non-Zero Wage Growth

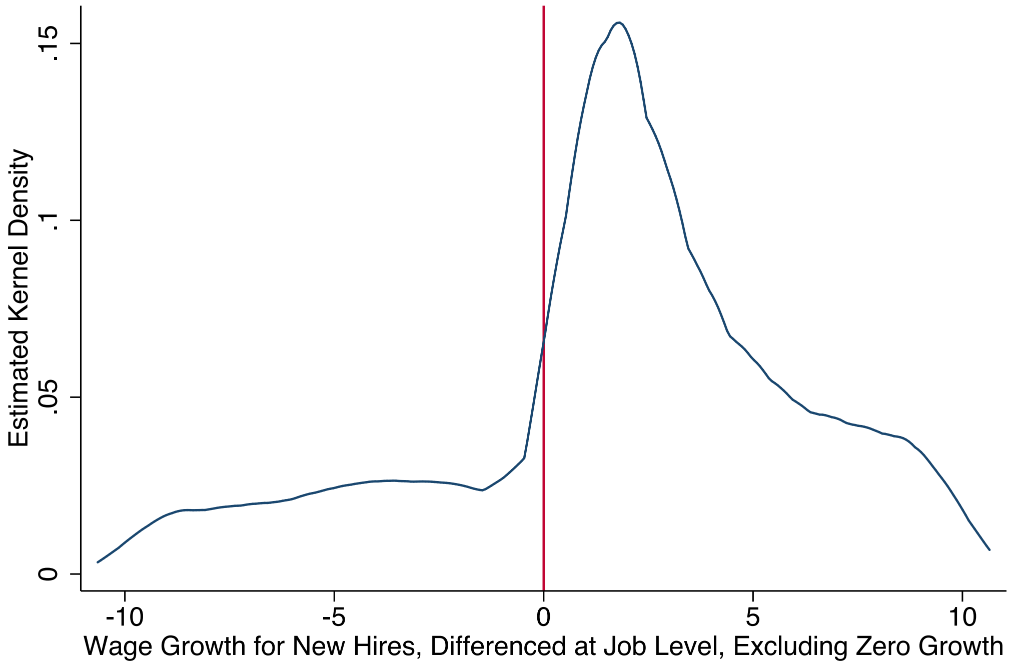

Notes: this graph is the distribution in the growth of wages for new hires, excluding zeros, from Burning Glass. A job is an establishment by job-title by salary type by pay frequency unit. Wages are averaged by job-quarter. Wage growth is the growth in wages between two consecutive vacancies posted by the same job. The wage growth distribution is truncated at ±10%. Kernel density estimation uses an Epanechnikov kernel with a bandwidth of 0.65. The McCrary test tests the null hypothesis that the density function of wage changes is continuous at zero.

are in Table 4. As expected, wages are more likely to rise than to fall. Table 4 shows that the finding is robust across several specifications, including after reweighting to target the occupational or regional distribution of jobs, or excluding low wage jobs—in order to strip out the effect of minimum wages. In Appendix Section B, we repeat the analysis at annual frequency, with similar results.

Still, wages for new hires do sometimes fall at the job-level. This result may limit the importance of the downward constraint, and echoes similar results from Elsby, Shin, and Solon (2016) and Jardim, Solon, and Vigdor (2019) amongst others.

3.4 Wage Increases Are Cyclically Sensitive, Wage Decreases Are Not

We now show that the probability of wage increase is sensitive to business cycles, while the probability of wage decrease is not. Again, this finding suggests a constraint on cutting wages between vacancies. Firms let wages respond to cyclical conditions by varying whether wages increase—while rarely lowering wages irrespective of labor market tightness.

We estimate time varying probabilities of the change, increase, and decrease in the wage

Figure 4: Probability of Job-Level Change in Wage for New Hires

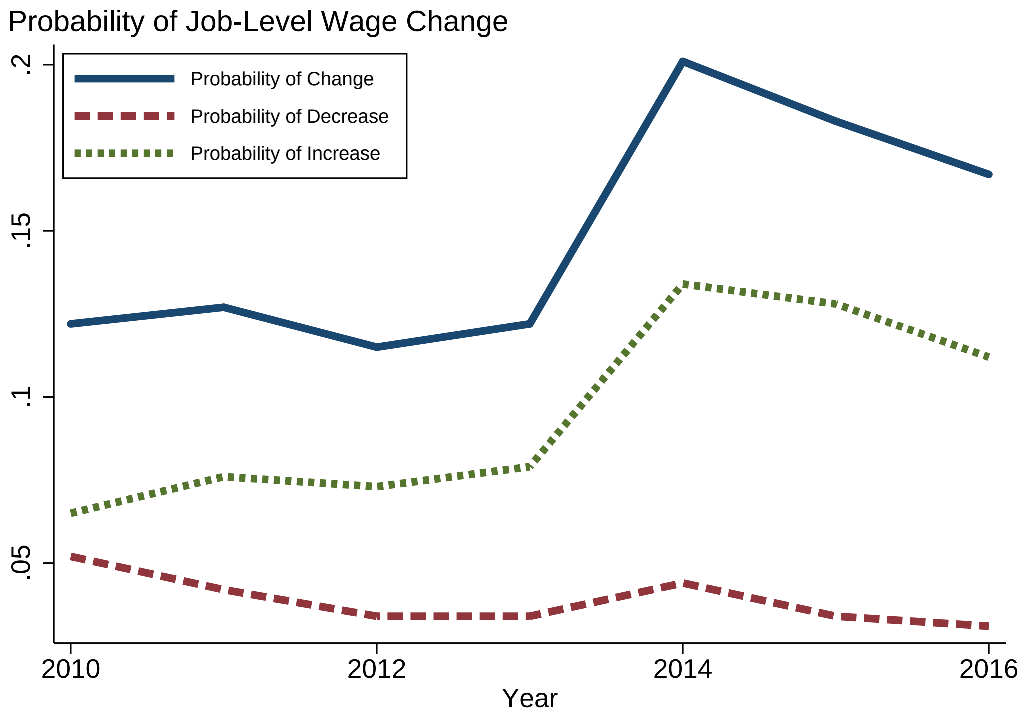

Notes: this graph estimates the job-level probability of wage change, increase and decrease, using the same method as in table 4, separately for each year, using Burning Glass data.

for new hires at the job level. We estimate these probabilities separately for each year of our sample, over 2010-2016, using the hazard model of subsection 3.2.

Figure 4 shows the results. As the labor market tightens over 2010-2016, the probability of wage change rises—as expected, given that wages rise over this period. However, the probability of wage change rises entirely because wage increases are more likely. Wage decreases are not more likely as the business cycle evolves. Thus wage increases for new hires are cyclically sensitive, and wage decreases are not.[19]

3.5 Wages For New Hires vs. Continuing Workers

Our finding, that the wage for new hires changes infrequently and falls rarely, is novel. We provide context with a fact that has previously been documented. Workers in continuing employ-

[19] Figure 4 uses variation from only one business cycle expansion. In further support, Appendix Table 5 shows that the probability of increase is more cyclical than the probability of decrease, with respect to state business cycles. We calculate the probability of wage increase and decrease within each state-quarter. We regress these probabilities on the growth in employment for each state-quarter. The probability of wage increase comoves strongly with employment growth, the probability of decrease does not. See Jo (2019) for a related finding with continuing workers' wages.

Figure 5: Duration of Unchanged Wage for Continuing Workers and New Hires

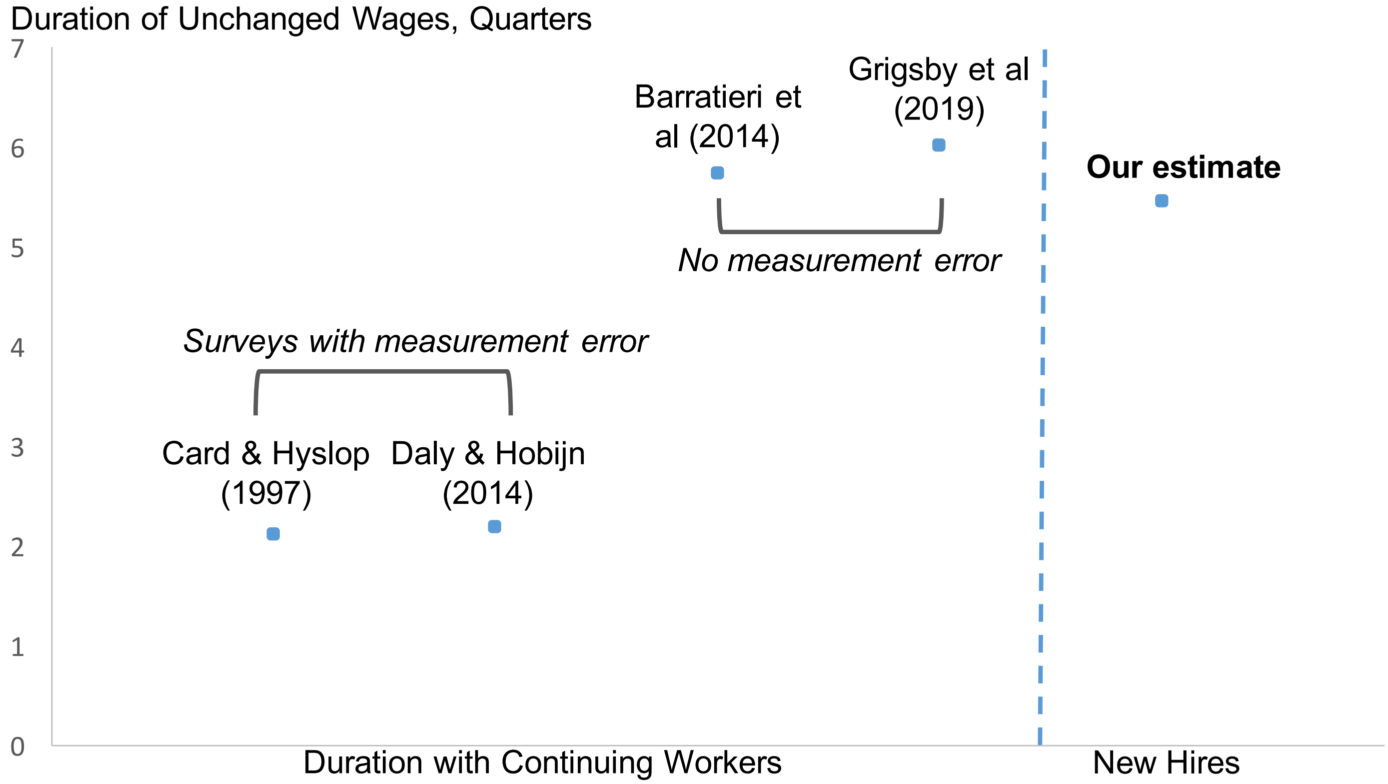

Notes: this graph plots the implied duration for which wages are unchanged from four papers that study continuing wages using payroll and survey data, alongside our estimate for new hires wages using Burning Glass data.

ment—as opposed to workers newly hired into jobs—rarely experience wage changes.

The duration for which wages do not change is similar, for new and continuing jobs. Figure 5 presents estimates of the duration that base wages are unchanged in continuing jobs. Two estimates are close to ours: the estimate of Grigsby, Hurst, and Yildirmaz (2018), which studies high quality payroll data; and the estimate of Barattieri, Basu, and Gottschalk (2014), which corrects for measurement error in survey wages.[20]

Our findings suggest that new and continuing wage changes are governed by the same underlying forces. Previous work conjectures that new and continuing wages behave similarly due to internal equity between new hires and continuing workers, or firm wide pay scales (Bewley, 2002). Then wages change infrequently across successive new hires. Our finding supports this argument.

Our finding that wage setting is similar for new and continuing jobs is not obvious. Some

[20] Other estimates are from survey data without correcting for measurement error, which biases downwards the estimated duration for which wages are unchanged.

plausible mechanisms predict the opposite pattern. As one example, implicit contracting models imply that continuing wages should be rigid downwards, while wages in new hires should be *flexible* downwards (Harris and Holmstrom, 1982; Beaudry and DiNardo, 1991).[21] 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, who do not have a reference point of their own past wage (Eliaz and Spiegler, 2014).

4 Wage Cyclicality

This section asks whether the wage for new hires responds differently to business cycle contractions and expansions at the job level. We have two key results. First, across successive vacancies posted by the same job, the wage for new hires does not fall during contractions, but does rise during expansions. Second, wage flexibility *upward* is state dependent, in a way that is consistent with downward rigidity.

4.1 Regional Unemployment Variation

In our regressions, we study the response of wages to unemployment, to measure wage cyclicality as in Bils (1985). We study regional business cycles, to avoid the problem of a relatively short time series in our data. State level unemployment is measured with noise. We instrument for state-level unemployment with an administrative measure of employment, from the Quarterly Census of Employment and Wages (QCEW), to avoid attenuation bias.

States are a natural definition of a regional labor market. Since 2010, interstate migration has been relatively low, and mostly unrelated to cyclical considerations (Yagan, 2016; Beraja, Hurst, and Ospina, 2016). Moreover there is substantial regional business cycle variation during this period. Various states (e.g. the District of Columbia and New York) saw rising unemployment during 2010-2012 due to the prolonged impact of the Great Recession. Other states saw rising unemployment due to the faltering labor market recovery in 2013 (e.g. Illinois, Oklahoma, Massachusetts and Ohio). A third group of states suffered in 2015-6 due to falling oil prices (e.g. North Dakota, Texas, Wyoming, New Mexico, Alaska and Oklahoma). Appendix Section A documents further statistics about regional business cycles over this period.

[21] Within jobs, risk neutral firms insure risk averse workers, by offering them downwardly rigid wage contracts. The wage for new hires, as firms and workers enter a new implicit contract, is not constrained by the insurance motive. Beaudry and DiNardo (1991) present evidence that continuing workers have more rigid wages than new hires, though their interpretation of the data is disputed (Hagedorn and Manovskii, 2013).

4.2 Benchmark Specification

Our benchmark regression for measuring wage cyclicality, at the job level, is

$$\Delta \log w_{jst} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{jst}. \tag{1}$$

This equation is similar to the standard regression of Bils (1985). w_{jst} is the nominal wage for a new hire in job j and quarter t. We difference wages between the successive quarters in which the job posts a vacancy, which may be a gap of more than a single quarter.[22] This step isolates job-level wage changes. ΔU_{st} is the change in quarterly state level unemployment. γ_t is a time fixed effect. β and δ measure the sensitivity of the wage for new hires to regional unemployment. A more negative number indicates greater sensitivity. If $\delta < 0$, then wages comove more with unemployment during expansions, that is, when $\Delta U_{st} < 0$. If $\beta = 0$, then wages do not comove with unemployment during contractions. We instrument for ΔU_{st} and $I\left[\Delta U_{st} < 0\right] \times \Delta U_{st}$ with $\Delta \log\left(\text{employment}_{st}\right)$ and $I\left[\Delta \log\left(\text{employment}_{st}\right) < 0\right] \times \Delta \log\left(\text{employment}_{st}\right)$, where $\Delta \log\left(\text{employment}_{st}\right)$ is state-quarter employment growth from the QCEW.[23]

We study the same sample as in section 3, that is, jobs that post multiple vacancies, averaging wages at the job-quarter level. Time fixed effects sweep away aggregate variation, to focus on regional variation. Time effects also control for variation in the national price level. Therefore our results measure real wage rigidity, deflated by national prices. For a valid structural interpretation, regression (1) must assume that unemployment is driven by labor demand shocks. We explore robustness to this assumption in our regressions.

4.3 Job-Level Wages Respond to Falls in Unemployment, But Not Rises

We turn to the first key empirical result of the section. Figure 1, previously shown in the introduction, illustrated that when unemployment rises, the wage for new hires does not fall—meanwhile wages do rise as unemployment falls.

Table 5 confirms these results by estimating regression equation (1). In Column (1) of Table 5, β is not significantly different from zero, and indeed is slightly positive—thus the wage for new hires does not fall during contractions. Meanwhile δ is negative and statistically significant. Wages are more sensitive to expansions than contractions in unemployment, and rise during expansions. The results—both that β is near zero and δ is significantly negative—are ro-

[22] Differencing wages across successive quarters allows an easy comparison to papers such as Hagedorn and Manovskii (2013) and Gertler, Huckfeldt, and Trigari (2016). In those papers, the regression differences across the successive quarters in which workers are hired into new jobs, analogous to our regression.

[23] Appendix Table 14 reports the first stage regression projecting quarterly state unemployment changes onto employment growth. As expected, the two series are closely correlated.

bust across several specifications. In column (2) we add in state-specific trends, and in column (3) we reweight to the occupational distribution of jobs in the US economy, to ensure representativeness. We reweight at the 6 digit SOC code level, using the 2014-2016 Occupational Employment Statistics.

Column (4) drops the $I[\Delta U_{st} < 0]\,\Delta U_{st}$ term from our benchmark regression (1), and instead measures the average sensitivity of wage growth to unemployment changes. On average wages do comove negatively and significantly with unemployment—but this average comovement is entirely driven by expansions and not contractions.

We doubt labor supply fluctuations could rationalize the sharp asymmetries that we document.[24] Nevertheless, column (5) studies an instrument for labor demand and again finds downward rigidity. We instrument for state unemployment using a Bartik-style instrument based on states' regional exposure to the global oil price.[25] Again δ is negative and significant and β is insignificantly different from zero. Thus wages do not fall during contractions, and are more rigid downward than upward, in response to labor demand shocks. In time series data, oil price shocks seem to be symmetric (Kilian and Vigfusson, 2011). So, the asymmetry in column (5) comes from the response of wages and not from asymmetry in the process for labor demand.

The identifying assumption for our instrument is that states who are exposed to contractions in the global oil price do not receive particularly large labor supply shocks at the same time. This assumption is similar to Acemoglu, Finkelstein, and Notowidigdo (2013) and Allcott and Keniston (2017). The assumption seems plausible in our setting. The variation in this instrument comes from the large contraction in the oil price in 2015 and the increase beforehand—Appendix Figure 5 displays the oil price. Appendix Figures 3 and 4 show that the regional contractions during this period come mostly from oil producing states, such as Texas, Wyoming and Alaska.

A large share of vacancies post while state unemployment is increasing, letting us estimate asymmetries despite the national labor market recovery over 2010-2016. Where does this variation come from? Appendix Figure 3 plots the share of vacancies in each state that experience rising state unemployment. Appendix Figure 4 plots the share of vacancies in each year that

[24] Still, there are important labor supply shocks that differentially affect states over this period, such as the extension and elapsing of unemployment benefit durations (e.g. Hagedorn, Karahan, Manovskii, and Mitman, 2013).

[25] The first stage regression is $\Delta U_{st} = \sum_s \left[\beta_s \Delta \log\left(\text{oil price}_{t-1}\right) + \gamma_s I\left(\Delta \log\left(\text{oil price}_{t-1}\right) < 0\right) \Delta \log\left(\text{oil price}_{t-1}\right)\right] + \text{error}_{st}$, where α_s, β_s and γ_s are estimated, similarly to Nakamura and Steinsson (2014). There are many instruments, which biases the estimates *towards* OLS, and therefore strengthens the interpretation of our finding, because our IV estimate of the downward wage rigidity coefficient δ is greater in magnitude than our OLS estimate. Nakamura and Steinsson report for their instrument that the standard error is unbiased, because of the high R^2 of the instruments as a whole. Though we cluster standard errors by state in other regressions, we cluster standard errors by both state and year in this regression, following the recommendation for inference in Bartik instruments by Adao, Kolesár, and Morales (2018).

experience rising state unemployment.

We now present various robustness tests concerning our finding of job-level wage rigidity.

4.3.1 Downward Wage Rigidity at the Job-Level—Robustness Tests

Table 6 groups together some robustness tests about our key finding that wages are rigid downward and flexible upward. 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. Row 1 is our baseline specification from column (1) of Table 5.

Rows 2 and 3 suggest that strategic vacancy posting by firms does not explain our results. The concern is that firms might initially attach low wages to their vacancies, and then gradually post vacancies with higher wages until the vacancy is filled. This behavior might lead to asymmetric wage movements. Row 2 excludes the vacancies that are posted in a quarter immediately after another vacancy of the same job. These are the vacancies for which such strategic behavior is most likely to matter. Row 3 studies the benchmark regression at annual frequency, by taking the mean wage across vacancies posted in each job and year. At lower frequency, strategic vacancy posting is less likely. In both cases the magnitude of the estimated coefficient increases, strengthening our result.

Rows 4 and and 5 suggest that selection bias does not explain our results. The concern is that wages only appear in the data when firms choose to post vacancies. There is a classic "missing data problem" that may lead to selection bias. In row 4, we control for the length of time that has elapsed since a job previously posted a vacancy. If more time has elapsed since a job last posted, then there is more missing data and a hence greater concern of selection bias. In row 5 we control for selection bias using a standard Heckman (1979) estimator.[26] In both cases, the degree of downward wage rigidity remains large and significant.

Row 6 reweights vacancies to target the regional distribution of employment in the Quarterly Census of Employment and Wages. Rows 7-8 seasonally adjust by either applying the Bureau of Labor Statistics' X-11 algorithm to unemployment, or adding state by quarter of year fixed effects to the regression. Row 9 studies only wages that post a point wage, instead of a wage range.[27] Row 10 removes wages with bonuses from the data. Row 11 studies a broader definition of a job. We consider a new definition of a job, as a job title within an establishment, while

[26]Specifically, we consider a version of the standard Heckman model, in which the probability that a job posts a vacancy in a given quarter depends on the level of unemployment and log employment in that quarter, as well as on the regressors in equation (1).

[27]Appendix Table 6 further investigates the issue of wage ranges. This table shows that the share of vacancies posting wage ranges instead of point wages does not vary significantly with regional unemployment. So, vacancies posting ranges and point wages do not have different cyclical properties.

pooling across pay frequencies.[28] Row 12 removes time fixed effects. In all cases, the coefficient is negative and significant, implying that the wage for new hires is more rigid downward than upward.

Our Appendix contains further robustness exercises. These exercises use different sources of variation to show that downward wage rigidity is pervasive. Appendix Table 12 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. Appendix Table 7 shows that wages are more rigid downward than upward at the 3 digit industry level. Appendix Table 8 shows the same result, using 2- and 3-digit industry by state variation. These regressions include state-by-time fixed effects, sweeping away state level labor supply shocks over this period, such as unemployment benefit extensions (Hagedorn, Karahan, Manovskii, and Mitman, 2013). Appendix Table 9 studies real wages for new hires at the city level, deflated by BLS measures of city prices. These regressions show that real wages are more rigid downward than upward, and that the magnitude of nominal and real wage rigidity is similar. Appendix Table 10 estimates our baseline regression for five broad occupations, and finds that downward wage rigidity is pervasive across all broad occupations. Appendix Table 11 estimates our baseline regression by quartiles of establishment size, and finds that downward wage rigidity is pervasive across establishment size. Appendix Table 11 also estimates our baseline regression across the four quartiles of the wage distribution, and again finds that downward wage rigidity is widespread.[29]

Now, we consider four additional issues that might undermine our finding of downward wage rigidity at the job level. Each is important but we suspect none fully overturn our results.

4.3.2 Worker Composition and Hiring Standards

First, our estimates control for job composition but they do not control for worker composition. Changing worker composition might offset wage rigidity at the job level. Suppose that wages do not fall during contractions, but firms are able to hire more productive workers into the job. Then firms' marginal costs are still flexible downward.

We investigate this concern and do not find that worker composition offsets wage rigidity at the job level. Our data has an advantage for examining worker composition: we can study changes in education requirements attached to the vacancy. Education is a measure of worker skill. So, education requirements proxy for changing worker composition. We study the cycli-

[28] In the baseline, a job is a job title within an establishment at a given pay frequency (e.g. hourly or annual).

[29] Table 13 decomposes wage growth at the job level its extensive margin and intensive margin components, and finds that both components contribute to wages being more rigid downward than upward.

cality of education requirements in Appendix Section C.2. We do not find any significant co-movement between education requirements and unemployment. So, changing education requirements do not seem to offset downward wage rigidity.

Education requirements are also a measure of hiring standards. So, our results suggest that changes in hiring standards do not offset downward wage rigidity at the job level.

We do not argue against the importance of worker composition for wage and unemployment dynamics. Rather, we think worker composition is unlikely to nullify wage rigidity at the job level.[30]

4.3.3 Posted Wages versus the Wage for New Hires

A second potential complication is that our data records the wage posted on vacancies. Suppose that workers bargain, so that there is a gap between wage posted in a vacancy and the wages for newly hired workers. If this gap varies over the business cycle, it could affect the interpretation of the job-level wage rigidity that we estimate.

We suspect this concern does not undermine our result for three reasons. First, we previously showed that wages posted in Burning Glass closely track other measures of the wage for new hires from survey and administrative data—which points against a gap between wages posted in Burning Glass, and the wage for new hires.

Second, in Appendix Section C.4, we ask whether the difference between measures of the wage for new hires in Burning Glass and measures of the wage for new hires from survey data correlate with unemployment changes. If so, then the gap between posted and realized wages for new hires might vary over the business cycle. However, we do not find that the difference correlates with unemployment changes.

Third, survey evidence suggests that the majority of workers accept wages posted on vacancies, which again minimizes the issue due to bargaining (Hall and Krueger, 2012).

4.3.4 Wage Growth after Being Hired and the User Cost of Labor

A third potential complication is that our dataset measures wages at the point of hiring. But job creation is a long term decision. So, the present value of wages paid to new hires may also matter for job creation (Kudlyak, 2014). This consideration might affect our baseline result. If firms reduce wage growth for workers after they are hired, they might offset downward wage rigidity at the point of hiring (Elsby, 2009). In the language of Kudlyak (2014), the “user cost of labor” could be more flexible than the wage for new hires.

[30] Mueller (2017), Grigsby (2019) and Carrillo-Tudela, Gartner, and Kaas (2020)show worker composition or hiring standards are important for unemployment and wage dynamics.

Appendix section C.3 directly investigates whether workers' wage growth after being hired can offset downward wage rigidity at the point of hiring. We cannot see wages after workers are hired in our dataset. Instead we study workers' wage growth after they are hired, relative to the wage when they are hired, from longitudinal worker-level survey data. Our method follows the state of the art from Kudlyak (2014) and Basu and House (2016).[31]

We do not find evidence that wage growth after workers are hired offsets downward wage rigidity at the point of hiring. Our results suggest the opposite. If workers are hired during a contraction, then their subsequent wage growth seems to be *higher*. So, the present value of wages does not seem to fall during contractions.

4.3.5 Establishment Variation

Establishments hire multiple types of jobs. Even though job level wages seem to be downwardly rigid, establishments might avoid wage rigidity by changing their mix of jobs. For example, consider a Starbucks establishment in which wages are downwardly rigid for "senior baristas" and "junior baristas". During expansions, Starbucks hires higher wage senior baristas. During contractions, Starbucks hires lower wage junior baristas. Either way, newly hired workers brew coffee. The wage for new hires falls despite downward rigidity at job level, without any effect on the output of the Starbucks. So, the establishment avoids wage rigidity by changing the mix of senior versus junior baristas.

We start with an establishment level version of wage cyclicality regression. We study the regression

$$\Delta \log w_{et} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta I \left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{et}. \quad (2)$$

w_{et} is the mean nominal establishment wage, pooling across all jobs posted by an establishment in a given quarter. We difference wages between the successive quarters in which the establishment posts a vacancy. This step isolates establishment-level wage changes. The establishment-level regression has a different outcome variable from our job-level regression that tests for downward wage rigidity, equation (1)—but otherwise, the two regressions are identical.

Table 7 reports the results. In all columns of Table 7, β is not significantly different from zero—thus the wage for new hires does not fall during contractions at the establishment level. Meanwhile δ is negative and statistically significant. At the establishment level, wages are more sensitive to expansions than contractions in unemployment, and rise during expansions. The magnitude of downward wage rigidity is similar at the job and establishment level. [32]

[31] In Appendix section C.3, we study wage growth workers have been hired into a given job, relative to the wage paid to workers when they are hired into the job. So, job composition does not affect this supplementary analysis.

[32] In Appendix Section C.5 we find no evidence that establishments alter the mix of jobs that they hire, in order to reduce the effects of downward wage rigidity.

4.4 Downward Rigidity and State Dependent Flexibility Upward

We turn to the second key result of this section pointing to downward wage rigidity. The wage for new hires displays state dependent flexibility *upward*. The form of state dependence is consistent with downward wage rigidity, and has not been documented in prior work.

Let us explain the prediction of downward wage rigidity for state dependence. Suppose that wages are downwardly rigid. A simple model for downward wage rigidity is

$$w_t = \max\left[w_{t-1}, w_t^*\right]$$

$$w_t^* = b + \phi y_t \quad b, \phi > 0.$$

Wages today, w_t, are the maximum of previous wages w_{t-1}, and a frictionless wage w_t^*. If the frictionless wage is low, wages today may be constrained by previous wages. The frictionless wage depends positively on labor demand y_t.[33]

In this simple model, wage flexibility upward is state dependent, due to downward wage rigidity. After a large contraction in labor demand at $t-1$, w_t is much greater than w_t^*. Then if a slight rise in labor demand follows at time $t+1$, we have $w_{t+1} = w_t > w_{t+1}^*$, that is, wages do not rise as labor demand marginally increases from the trough of the contraction. Wages are "trapped too high" by the downward constraint. Suppose instead that the economy has been expanding, so downward constraints do not bind and $w_t = w_t^*$. Then after a rise in labor demand, $w_{t+1} = w_{t+1}^* > w_t$, and wages rise after the increase in labor demand.

Figure 6 shows state dependence consistent with downward rigidity. On the y-axis is wage growth between two consecutive vacancies for the same job, from Burning Glass. On the x-axis is the growth in quarterly state level unemployment between the quarters in which the vacancies are posted, from the Bureau of Labor Statistics. We plot the data separately for state-quarter observations in which there has been a contraction in the past (circles), and observations in which there has been an expansion in the past (triangles). There has been a contraction if quarterly state unemployment rose over the previous three years, otherwise there has been an expansion. Figure 6 shows state dependence consistent with downward rigidity. When there has been a contraction in the past, then falls in unemployment do not increase wages. When there has been an expansion in the past, then falls in unemployment do increase wages. Due to downward wage rigidity, wages do not respond to increases in unemployment, neither when there has been a contraction nor an expansion in the past.

[33] This simple model is similar to, amongst others, Schmitt-Grohé and Uribe (2016), Chodorow-Reich and Wieland (2017) and Dupraz, Nakamura, and Steinsson (2016).

Figure 6: Wages for New Hires and Unemployment—Contractions Versus Expansions in Past

Notes: the graph plots binned wage growth for new hires, from Burning Glass, and binned state by quarter unemployment changes, from the Bureau of Labor Statistics. To construct wage growth, we take the mean wage within each job and quarter, and then take log differences at the job level. We use 50 bins, partial out time fixed effects, and add a non-parametric regression line. The triangles plot data when there has been an expansion in the past. The circles plot data when there has been a contraction in the past. We define a contraction as a period in which state unemployment rose over the previous three years, otherwise there is an expansion.

We confirm this finding with regressions. We estimate the regression

$$\Delta \log w_{jst} = \alpha + \gamma_t + \kappa \Delta U_{st} + \nu \Delta U_{st} \times I\left(U_{s,t-1} - U_{s,t-13} < 0\right) + \text{controls}_{jst} + \varepsilon_{jst}.$$

The dependent variable is quarterly job-level wage growth for new hires, from Burning Glass. The independent variable is the change in state-quarter unemployment. We interact state-quarter unemployment changes with an indicator for whether state unemployment fell over the previous three years. As before, we project unemployment changes on employment growth from the Quarterly Census of Employment and Wages to deal with measurement error.

We restrict the sample only to observations for which $\Delta U_{st} < 0$ in order to study wage flexibility upward. Therefore κ measures the sensitivity of the job-level wage for new hires to falls in state unemployment, when unemployment has contracted over the previous three years. If κ is near to zero, then wages grow little as unemployment falls, in the aftermath of a previous contraction. If v is significantly negative, then wages are more sensitive to falls in unemployment, in the aftermath of an expansion over the previous three years. So if v is negative, there is state dependent wage flexibility upwards: wages are more sensitive to increases in labor demand when the economy has previously been expanding.

Table 8 presents the results. Across all specifications, v is large in magnitude and significantly negative. Therefore in the aftermath of expansions, the wage for new hires is more responsive to increases in labor demand. Wage flexibility upward is state dependent. Importantly, in the final column, we control for the level of unemployment. In standard models of wage setting such as Nash bargaining, the level of economic activity also affects wage flexibility. After adding this control, our estimate v of state dependence changes little.

This form of state dependence predicts wages should have been inflexible upward in 2010, in the aftermath of the Great Recession. Wages should have become progressively more flexible upward over the course of the recovery. We find evidence for precisely this phenomenon, in Appendix Section C.7. So, our estimates may be able to explain the puzzle of "missing wage growth" during the early recovery from the Great Recession. In 2014, the rapid decline in unemployment and slow growth of wages was deemed puzzling by many observers (e.g. Federal Reserve Bank of Atlanta, 2014). But state dependence due to downward rigidity may have contributed to this pattern.

4.5 Reallocation Between Establishments and Downward Wage Rigidity

In this section, we showed downward wage rigidity at the job and establishment level. We now consider whether reallocation between establishments might undo the effects of establishment level wage rigidity. Let us explain the concern with an example. Suppose that, on average, wages are downwardly rigid at the Starbucks establishment, but there is a neighboring establishment of Dunkin'. On average, wages are higher at Starbucks than Dunkin'. After a contraction in labor demand, Starbucks stops hiring. However Dunkin', with its lower wages, is able to hire the workers who cannot find jobs at Starbucks. Either way, the same workers still make coffee. More generally, reallocation of workers between establishments might undo downward wage rigidity at the establishment level.

We test for the concern by asking whether the share of low wage jobs in the overall labor market increases during contractions. For each state and quarter, we calculate the share of high wage vacancies, that is, vacancies with above median wages in Burning Glass. We regress the quarterly change in the high wage state share of vacancies, on the change in quarterly state unemployment. The regression is identical to regression (1), except for the outcome variable—which is the quarterly change in the state share of high wage vacancies.

Table 9 presents the results. Row (1) of column (1) shows that when state unemployment rises by one percentage point, the share of high wage jobs falls by a statistically insignificant 0.6 percentage points. Moreover, row (2) of column (1) shows that the share of high wage vacancies at the state level does not respond significantly differently to rises versus falls in unemployment. Thus the state share of vacancies are not moving in a way that offsets the asymmetric response of wages to contractions versus expansions. Columns (2) and (3) study the same regression after adding in state trends and reweighting to target regional employment. The regression coefficients are noisy and unstable, but none of them suggest that the state share of high wage vacancies responds differently to rises versus falls in unemployment.

Equally, our evidence does suggest that the high wage share of vacancies falls slightly during recessions. Column (4) of Panel A reports the coefficient from regressing the quarterly change in high wage vacancies on the quarterly change in state unemployment. This regression studies the average effect, and does not separate out the effect of expansions versus contractions in unemployment. On average, as unemployment rises, the high wage share of vacancies falls very slightly. However, the estimated coefficient is small. This finding is somewhat consistent with previous work, such as Barlevy (2002), which finds that workers often switch to lower wage jobs during recessions. The small effect size that we estimate may reflect other factors such as upskilling (Hershbein and Kahn, 2016), which raises the share of high wage vacancies during recessions.

5 Job Composition and the Average Wage for New Hires

To summarize the results so far: we provided new evidence that the wage for new hires is rigid downward and flexible upward, at the job and establishment level. No previous work detects this asymmetry. But previous work studies the *average* wage for new hires, from worker level survey data that controls for worker characteristics and averages across jobs (Haefke, Sonntag, and Van Rens, 2013; Basu and House, 2016). This section shows that due to job composition, average wages may have higher variance than job or establishment wages. So, regressions with average wages have limited power to detect downward wage rigidity for new hires. We also explain a related issue: job composition might create a form of omitted variable bias (Hagedorn and Manovskii, 2013, Gertler, Huckfeldt, and Trigari, 2016). In our data, this latter issue is less important.

First, we precisely define a job level measure of the wage for new hires, to contrast with the average wage for new hires used in prior work. Consider an economy with I job types, S states, and T time periods. 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 v_{ist}.

Our dataset measures growth in the *job-level* wage for new hires, $\Delta \log w_{ist}$. That is, we observe growth in the wage for new hires, for the same job, in the same state, between successive quarters.[34] Previous researchers measure the *average* wage for newly hired workers from survey data without information on jobs or establishments, such as the Current Population Survey or the National Longitudinal Survey of Youth. Researchers then calculate the average log wage of newly hired workers, and approximate the growth in the average wage for new hires by

$$\Delta \overline{\log w_{st}} = \sum_i \nu_{ist} \log w_{ist}, \tag{3}$$

which is the change in average log wages.[35]

Previous work often measures the average wage for new hires after controlling for worker characteristics. There are two methods. First, several papers regress the wage of newly hired workers on worker characteristics or worker fixed effects. The mean of the residual is the average wage for new hires controlling for worker characteristics (e.g. Haefke, Sonntag, and Van Rens, 2013; Basu and House, 2016). Second, several papers run regressions directly on worker level data on the wage for new hires, and add worker characteristics or worker fixed effects to the regression (e.g. Gertler, Huckfeldt, and Trigari, 2016). The first and second procedure are numerically equivalent. But importantly, while both methods construct the average wage controlling for worker characteristics, neither fully controls for job characteristics.

Estimates using average wages generally 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 equation (3) yields

$$\underbrace{\Delta \overline{\log w_{st}}}_{\substack{\text{average} \\ \text{wage growth}}} \approx \overbrace{\sum_i \nu_{ist} \Delta \log w_{ist}}^{\substack{\text{job level} \\ \text{wage growth}}} + \underbrace{\sum_i \log w_{ist} \Delta \nu_{ist}}_{\substack{\text{wage growth due} \\ \text{to composition}}}. \tag{4}$$

Average wage growth depends on two components: job-level wage growth, and wage growth due to composition. 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 (4) 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 (4) falls, so average

[34] Recall that in our empirics, we define a "job" as a job title by establishment unit.

[35] For practical reasons, researchers typically study the change in the average of log wages, instead of the change in the log of average wages. Equation (3) has an s subscript for comparability with what follows, but researchers normally study average wages for new hires at the national and not the state level.

wages fall.

From inspecting equation (4), 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 (4) has low variance. But suppose that the share of any given job type, v_{ist}, is volatile. Then the second term on the right hand side is large. So, average wage growth has high variance even if job-level wage growth has low variance. 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 many baristas. Then average wages will vary, even if wages change little for either bankers or baristas.

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 (4), average wages must *systematically* fall during recessions, even if job-level wages do not fall. Regression estimates from average wages could suffer from omitted variable bias, as emphasized by Solon, Barsky, and Parker (1994), Hagedorn and Manovskii (2013) and Gertler, Huckfeldt, and Trigari (2016). But in subsection 4.5, we showed that unemployment changes do not seem strongly correlated with job composition over our sample period, and job composition does not display the sharp asymmetries apparent in job-level wages. So, we do not focus on omitted variable bias due to job composition in this section.

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

$$\Delta \log w_{ist} = \alpha + \gamma_t + \beta_{\text{Job Level}} \Delta U_{st} + \delta_{\text{Job Level}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{ist}, \tag{5}$$

where ε_{ist} has bounded variance σ^2_{ist}. We are interested in $V\left[\hat{\delta}_{\text{Benchmark}}\right]$, the variance of the OLS estimator of $\delta_{\text{Benchmark}}$. If $\delta_{\text{Benchmark}}$ 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

$$\Delta \overline{\log w_{st}} = \bar{\alpha} + \bar{\gamma}_t + \beta_{\text{Average}} \Delta U_{st} + \delta_{\text{Average}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \bar{\varepsilon}_{st}. \tag{6}$$

This regression is the analogue of our job-level regression, with average wages as the outcome variable. If estimates of δ_{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 δ_{Average}, which we term $V\left[\hat{\delta}_{\text{Average}}\right]$, will be large.

In the following proposition, we show that job composition inflates the variance of $\hat{\delta}_{\text{Average}}$

relative to $\hat{\delta}_{\text{Benchmark}}$. Thus regressions using average wages may lack the power to detect downward rigidity, even if it is present at the job level.

Proposition 1. *For $S, T < \infty$, and if $\sum_i \log w_{ist} \Delta v_{ist}$ and $\sum_i \log w_{ist} \Delta v_{ist}$ are independent conditional on ΔU_{st}, then*

$$V\left[\hat{\delta}_{Average} | \Delta U_{st}\right] > V\left[\hat{\delta}_{Job\ Level} | \Delta U_{st}\right] \quad and \quad V\left[\hat{\beta}_{Average} | \Delta U_{st}\right] > V\left[\hat{\beta}_{Job\ Level} | \Delta U_{st}\right]$$

We collect the proof of Proposition 1, and all our other propositions, in Appendix Section D. The proposition makes a simple point. From equation (3), 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.[36]

In real-world data, job composition makes estimates that use average wages imprecise. We estimate $\hat{\delta}_{\text{Average}}$ in equation (6). For the outcome variable, we construct average wage growth for new hires at the state level, from Burning Glass, and from the Current Population Survey. To construct the average wage for new hires in the CPS, we follow the state-of-the-art procedure in Haefke, Sonntag, and Van Rens (2013), by measuring the average wage for new hires controlling for worker-level demographics and industry.[37] We study quarter-by-state data for 2010-2016, as in our benchmark regression. We report the standard error of $\hat{\delta}_{\text{Average}}$. We contrast with the standard error of $\hat{\delta}_{\text{Benchmark}}$. 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 $\hat{\delta}_{\text{Benchmark}}$, given that the regressor ΔU_{st} varies at the state level.[38]

Figure 7 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. Job composition does, indeed, inflate the variance of estimators of downward rigidity. The top row of Figure 7 reports the standard error of our job level estimate of downward wage rigidity, $\hat{\delta}_{\text{Benchmark}}$. The second row reports the standard error of $\hat{\delta}_{\text{Average}}$, the estimate of downward rigidity from average wages, with average wages for new hires from Burning Glass. The third row reports the standard error of $\hat{\delta}_{\text{Average}}$, with average wages for new hires from the Current Population Survey. The fourth row estimates $\hat{\delta}_{\text{Average}}$ using national wage

[36] Proposition 1 supposes that job composition $\sum_i \log w_{ist} \Delta v_{ist}$ and job level wage growth $\sum_i \log w_{ist} \Delta v_{ist}$ are independent, consistent with the evidence of subsection 4.5.

[37] We follow exactly the procedure Haefke, Sonntag, and Van Rens (2013), except that we also control for Census industry codes. To construct the average wage for new hires in Burning Glass, we follow the same procedure, and regress log wages for new hires on a set of state-by-quarter dummies, as well as dummy variables for the pay frequency and salary type of the wage. The state-by-quarter dummies are then mean log wages in the state-quarter.

[38] Clustering at the state level follows econometric best practice (Abadie, Athey, Imbens, and Wooldridge, 2017). See Solon, Barsky, and Parker (1994) for alternative ways to construct standard errors.

Figure 7: Estimates of Downward Rigidity in Job Level and Average Wage for New Hires

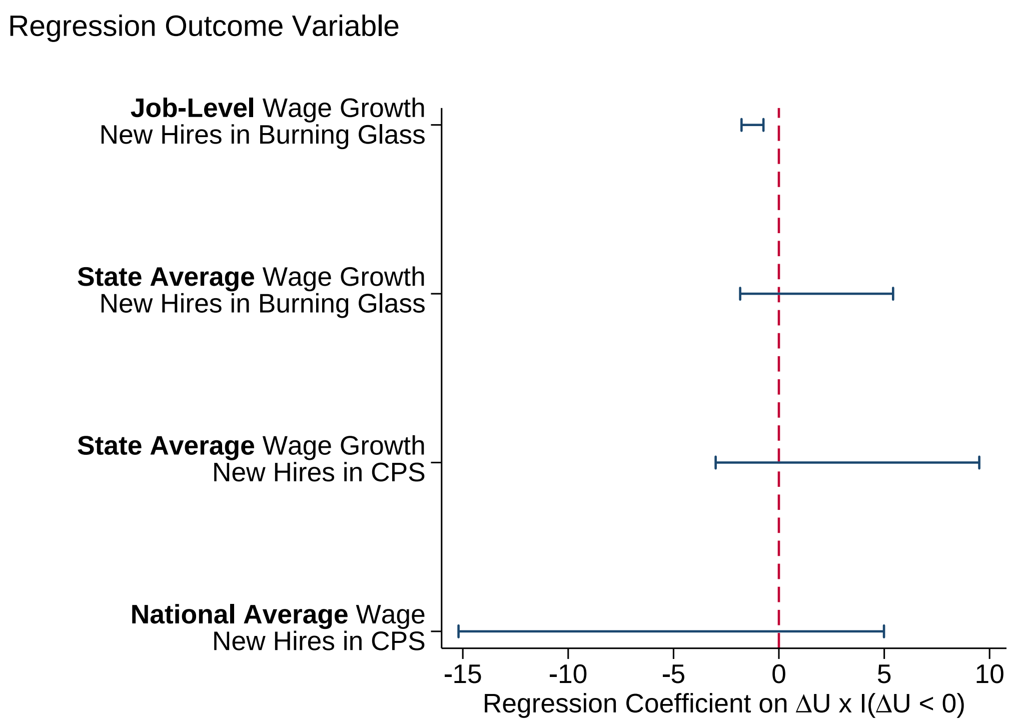

The top row reports the estimate of δ, from regression (5), which estimates downward rigidity with job-level data on the wage for new hires from Burning Glass. The next three rows report various estimates of δ from regression (6), which estimates downward rigidity with average wages for new hires from 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 Current Population Survey, and national average wages from the Current Population Survey. See Table 10 for details.

growth for new hires and national unemployment changes, for 1985-2006. The sample period and measure of wages is the same as Haefke, Sonntag, 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. Table 10 reports the point estimates and standard errors from the regressions in Figure 7. Columns (1), (3) and (5) report the specifications from row (2), (3) and (4) of Figure 7, respectively.

The reason why the job-level estimates are more precise is *not* because Burning Glass has a larger sample than the Current Population Survey. Figure 7 shows that estimates using average wages are equally imprecise in both Burning Glass and the CPS. The regressor varies at the state level. So, the degrees of freedom in the regression that studies job level wages is the number of states, and not the number of jobs.[39] Put differently, recall that we cluster standard errors

[39]To see this point formally, note that we can estimate $\hat{\delta}_{\text{Benchmark}}$ using *only* data aggregated at the state by

in the benchmark regression at the state level, given that the regressor varies at the state level. So, holding fixed the number of states, more observations within the state do not lower the standard error much.[40]

Table 10 also considers specifications that control for further information, but still do not isolate job-level variation. In all cases, estimates of downward wage rigidity remain imprecise. Evidently, job-level variation is essential to obtaining precise estimates. Column (2) studies average wages from Burning Glass and controls for detailed occupation and industry information. Column (4) studies average wages from the Current Population Survey data and controls for detailed occupation information.[41] Column (6) studies average wages for new hires from the National Longitudinal Survey of Youth and controls for worker fixed effects, following the procedure for measuring the average wage for new hires in Hagedorn and Manovskii (2013), Kudlyak (2014), and Basu and House (2016). This column also controls for the proxy for job composition developed by Hagedorn and Manovskii (2013).[42] In all cases, the confidence intervals are wide.

6 Job Level Wages and Unemployment Fluctuations

We now develop a model. In the empirics, we measured wage rigidity for new hires at the job level. The model shows that measuring job level wages is important, because this wage variation governs unemployment fluctuations. We study a simple extension of the standard Diamond-Mortensen-Pissarides labor search model with multiple job types, in order to have a precise definition of job level wages. We use the model to make two points. First, we show analytically that changes in the wage for new hires, at the job level, are particularly important for unemployment fluctuations. Second, we show that the wage rigidity in the data leads to large and asymmetric unemployment fluctuations.

---

quarter level. The benchmark regression equation (5) is numerically equivalent to the regression

$$\sum_i \nu_{ist} \Delta \log w_{ist} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta_{\text{Benchmark}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \sum_i \nu_{ist} \varepsilon_{ist}.$$

The left hand side variable is the state average of job-level wage growth within each quarter. Thus all variables in this regression vary only at the state-by-quarter level.

[40] Some papers studying wage rigidity for new hires consider regressions with worker-level panel data, of the form $w_{it} = \alpha + \beta U_t + \text{control}_{it} + \varepsilon_{it}$ where w_{it} is the wage for a newly hired worker and U_t is quarterly national unemployment. Some papers cluster standard errors at the worker level (Gertler, Huckfeldt, and Trigari, 2016). However, the regressor varies at the quarter level. So, clustering standard errors at the worker level understates the true standard error of the estimator, since there are many more workers than there are quarters. See Solon, Barsky, and Parker (1994), p. 13, for a lucid discussion of this issue.

[41] Note, however, that consistent occupation codes are only available in the CPS from 2012 onward.

[42] The wage measure is from Basu and House (2016). The NLSY is a long panel, hence one can construct a measure of the average wage for new hires that controls for worker fixed effects.

6.1 Model Setup

We study a standard Diamond-Mortensen-Pissarides model. Time is discrete and infinite. Unemployment fluctuations are driven by output per worker y_t, which follows an exogenous AR(1) process with mean value 1, that is

$$y_t = (1-\rho) + \rho y_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma^2). \tag{7}$$

y_t is a measure of labor demand.[43] There is a unit measure of homogeneous workers, who are either employed and producing output y_t, or unemployed and searching for work.[44] Workers are risk neutral, and derive utility from consumption only. Workers have discount factor $\beta \in (0,1)$ over future utility flows. Workers consume their wage in the periods that they are employed, and derive no flow utility from unemployment.

6.1.1 Wage Setting: Downward Rigidity and Job Types

We now introduce our two modifications to the standard DMP model: high and low wage job types, and downward wage rigidity at the job level.

Workers search for employment in either high or low wage job types during period t. In each job type, risk neutral firms post vacancies to hire with workers.

When a worker and firm hire at time t, wages are set. w_{it} is the real wage for a worker newly matched with a job of type i, and wages are fixed for the duration of the match. The wage for new hires in each job type satisfies

$$w_{Ht} = \max\left[w_{H,t-1}, \phi_H y_t^{\gamma}\right] \tag{8}$$

$$w_{Lt} = \max\left[w_{L,t-1}, \phi_L y_t^{\gamma}\right] \tag{9}$$

where

$$0 < \phi_L < \phi_H < 1.$$

This specification of wage setting has two implications. First, there are high and low wage job types. Since $\phi_H > \phi_L$, the wage for new hires is higher for job type H.

Second, the wage for new hires is more rigid downwards than upwards at the job level. Wages cannot fall between successive hires made by the same job type. Yet wages can rise if

[43] In practice, many shocks other than labor productivity may affect labor demand, such as monetary or fiscal policy. Labor productivity stands in for this more general set of shocks.

[44] We focus on job heterogeneity by assuming homogeneous workers. In practice, worker heterogeneity also matters for wage dynamics and unemployment fluctuations—see Solon, Barsky, and Parker (1994), Basu and House (2016) and Mueller (2017), amongst others.

labor demand y_t increases sufficiently, with a pass through from y_t into w_{it} governed by the parameters ϕ_i and γ. When wages rise, their dynamics are similar to wages with Nash bargaining, as in the canonical DMP model of Shimer (2005) and others.[45]

The non-parametric evidence of Figure 1 motivates our approach to modelling wages. In the data, the relationship between wage growth and unemployment changes is roughly piecewise linear, with a break around unemployment changes of zero. When unemployment falls, during expansions, the wage for new hires rises at the job level. However the wage for new hires does not fall during contractions, at the job level. Equations (8) and (9) also imply that wages are more sensitive to expansions in labor demand when the level of labor demand is high, in line with the finding of subsection 4.4.[46]

We do not model the underlying forces that might cause ϕ_H and ϕ_L to differ, nor that might generate downward wage rigidity. Hall (2005b) shows that in DMP models, a range of wages are consistent with bilaterally efficient bargaining between workers and firms. The wage setting in equations (8) and (9) are consistent with bilateral efficiency in a neighborhood of the steady state.

Since workers and firms are both risk neutral, and the job is a long term contract, what matters to both parties is the present value of wages. The timing of wage payments that we choose, with wages are fixed throughout the match, is a convenient normalization that emphasizes the role of the wage for new hires.[47]

6.1.2 Frictional Labor Market

There is a separate frictional labor market for each job type. Each labor market resembles that of the standard labor search model.

We model worker transitions between labor markets in a simple way. At the end of period $t-1$, an exogenous share ω of workers in each job type switch to being unemployed and searching for work in the other job type. The probability that a worker switches job types does not depend on whether she is employed or unemployed at the end of period $t-1$.[48] Also at the end of period $t-1$, an additional share s of the $l_{i,t-1}$ workers employed in job type $i = H, F$ separate from their jobs, in order to search for jobs of the same type.

[45] See Michaillat (2012), Dupraz, Nakamura, and Steinsson (2016) and Chodorow-Reich and Wieland (2017) for similar approaches to modelling downwards wage rigidity.

[46] Our model for wages also features positive pass through from labor productivity into wages, but no pass through from workers' outside option into wages (Jäger, Schoefer, Young, and Zweimüller, 2018).

[47] Moreover, the large literature on wage rigidity in continuing jobs shows that wages after the start of the hire are cyclically unresponsive, in line with our assumption here. We discussed new empirical evidence in favor of this assumption in subsection 4.3.4.

[48] Thus there is no on-the-job search—when workers switch job types, they first leave their current job, and then search for a new type of job.

Thus at the beginning of period t, the number of unemployed workers searching for jobs of type i satisfies

$$u_{it} = \frac{1}{2} - (1-\omega)(1-s)\, l_{i,t-1}, \tag{10}$$

since there is a measure 1/2 of workers either employed or searching for work in each job type at the start of period t, and $(1-\omega)(1-s)\, l_{i,t-1}$ workers remain employed from the previous period. Aggregate unemployment is $u_t = u_{Ht} + u_{Lt}$.

There is a large measure of risk neutral firms of each job type, with discount factor $\beta \in (0,1)$. Firms in each job type post v_{it} vacancies in total, to hire with the unemployed workers. In period t, total hires n_{it} are given by a matching function $n_{it} = M(u_{it}, v_{it}) = \Psi u_{it}^{\alpha} v_{it}^{1-\alpha}, \alpha \in (0,1)$. The key state variable governing each labor market is labor market tightness

$$\theta_{it} \equiv v_{it}/u_{it}. \tag{11}$$

The per-period cost of posting vacancies is $c > 0$. Vacancy posting costs capture firms' recruiting expenses, as they search for workers in the frictional labor market. The vacancy filling rate is $q(\theta_{it}) = \Psi \theta_{it}^{-\alpha}$. The vacancy filling rate is decreasing in θ_{it}—in a tight labor market, firms cannot find workers easily. Workers start working in the same period that they are hired.

If a worker finds a job in period t, they start producing output in the same period. The job finding rate of a worker searching for either job type is $f(\theta_{it}) = \Psi \theta_{it}^{1-\alpha}$. The job finding rate is increasing in θ_{it}—in a tight labor market, workers find jobs easily.

Tightness and employment comove positively. When the labor market is tight, firms hire many workers and employment rises. In particular, employment during period t satisfies

$$l_{it} = \frac{1}{2} - \left(1 - f(\theta_{it})\right)\left(\frac{1}{2} - (1-\omega)(1-s)\, l_{i,t-1}\right). \tag{12}$$

6.1.3 Firm Profits

Our model of the firm also follows the standard DMP setup. If a hire is filled at time t, it immediately starts to produce output. For periods $t+j$ in which a hire is not destroyed, the hire in job type i produces output y_{t+j}, common across job types, and pays job-type-specific wage w_{it} to the worker. The firm receives flow profit $y_{t+j} - w_{it}$.

The value of an unfilled vacancy depends on the chance that a vacancy is filled, and the cost of posting vacancies, as well as its continuation value. Then if K_{it} is the value of an unfilled vacancy and $J_{i,t,t}$ is the value in period t of a vacancy that is filled in period t, K_{it} is given by

$$K_{it} = -c + q(\theta_{it}) J_{it,t} + \beta\left(1 - q(\theta_{it})\right) \mathbb{E}_t K_{i,t+1}. \tag{13}$$

The value of a filled vacancy to a firm is the flow profit, and the continuation value, after deducting the risk of job destruction. $J_{i,t,t+j}$ is given by

$$J_{i,t,t+j} = y_{t+j} - w_{it} + \beta\left[(1-s)(1-\omega)\mathbb{E}_{t+j} J_{i,t,t+j+1} + [1-(1-s)(1-\omega)]\mathbb{E}_{t+j} K_{i,t+j+1}\right] \tag{14}$$

where $\mathbb{E}_{t+j}$ denotes the expectation conditional on time $t+j$ information.

6.1.4 Free Entry and Equilibrium

There is free entry in vacancy posting. Vacancy posting continues until the labor market becomes tight. Then vacancies are hard to fill, driving the ex ante value of vacancies to zero. Free entry implies

$$K_{it} \geq 0 \quad v_{it} \geq 0 \tag{15}$$

for all t with complementary slackness. When labor productivity rises, job creation becomes more profitable. Firms create many vacancies and the labor market tightens.

Our definition equilibrium is the same as the standard labor search model, other than the addition of multiple job types. An equilibrium is a collection of stochastic processes

$$\{l_{it}, v_{it}, \theta_{it}, u_{it}, w_{it}\}_{t=0}^{\infty}$$

for $i = H, L$, that satisfy the law of motion for unemployment (10), the definition of labor market tightness (11), wage setting equations (8) and (9), the Bellman equations for the value of an unfilled vacancy (13) and the value of a filled vacancy (14), and the free entry condition (15). The equilibrium is conditional on initial employment $l_{i,-1}$ for each job type and the AR(1) process (7) for y_t.

6.1.5 Job-Level and Average Wage Changes

Our modelling of job-level wages is the point of departure from the standard Diamond-Mortensen-Pissarides model. The standard model has a single job type—there is no distinction between job-level and average wages.

But our model admits a distinction between job level and average wage changes. Job-level changes in the wage for new hires are Δw_{Ht} and Δw_{Lt}. These are wage changes between hires of the same job type, in successive periods. Given downward rigidity, job-level wage changes may be positive, but are never negative. Average wages for new hires are

$$w_t^{\text{average}} = \frac{n_{Ht} w_{Ht} + n_{Lt} w_{Lt}}{n_{Ht} + n_{Lt}}.$$

Average wages aggregate across the high and low wage job types.

6.2 Job-level Wages Are Key for Unemployment Fluctuations

We now show the first main result coming from the model. Job-level wage changes are allocative for unemployment fluctuations. We derive a formula linking unemployment changes to wage changes, and show that *job-level* wage changes are what matter. So according to the model, we have measured a particularly important object in our empirics.

Proposition 2. *In a neighborhood of the steady state and to a first order*

$$\frac{\Delta \log u_t}{\Delta \log y_t} = -A + B \overbrace{\left[\frac{\mu \Delta \log w_{Ht} + (1-\mu)\Delta \log w_{Lt}}{\Delta \log y_t}\right]}^{\textit{average job-level wage growth}} \tag{16}$$

where $A, B > 0, \mu \in (0,1)$ and $\Delta x_t \equiv x_t - x_{t-1}$ is the difference operator, for constants A, B and μ defined in Appendix Section D.2.[49]

The left hand side of equation (16) is the response of aggregate unemployment in the economy to labor demand shocks y_t. The term in the square brackets of the RHS is the response of a weighted average of job-level wage growth to labor demand. $\Delta \log w_{it}$ is wage growth across successive hires in job type i, which depends on the wage setting equations (8) and (9). A and B capture other time-invariant factors affecting the sensitivity, such as the average profit on newly created jobs, and the elasticity of the matching function.

Equation (16) reveals two key insights. First, job level wage changes are allocative for unemployment fluctuations. The response of unemployment to labor demand shocks depends entirely on how job-level wages respond to labor demand shocks. When job-level wages are more flexible, so $\Delta \log w_{it}/\Delta \log y_t$ is higher, then unemployment is less sensitive to labor demand, and $\Delta \log u_t/\Delta \log y_t$ is smaller in magnitude.

Intuitively, in labor search models, job creation depends on *firms'* incentives to create new jobs. In turn, job creation governs unemployment fluctuations in labor search models. Firms' incentives to create jobs depend on how profitable the new jobs are. These profits, in turn, depend on job-level changes in the wage for new hires, which deduct from profits in new jobs.

Crucially, it is job-level and not average wage changes which matter. From equation (16), changes in average wages for new hires play no independent role in governing unemployment

[49]To derive this equation, we also use the approximation that, inflows into and outflows from employment are equal. This approximation is accurate at quarterly frequency (e.g. Shimer, 2005). In our first order approximation, we exploit that all equilibrium equations are continuously differentiable except the wage setting equations (8) and (9). The proof of this, and all other propositions, are also in Appendix Section D.

fluctuations. Regardless of how average wages change, a weighted average of job-level wage growth $\mu \Delta \log w_{Ht} + (1-\mu) \Delta \log w_{Lt}$ pins down unemployment fluctuations.

The second insight is that equation (16) reaffirms the basic argument of Pissarides (2009). It is the wage for new hires, as opposed to continuing jobs, which is key for unemployment fluctuations. The wage for new hires governs firms' profits on the new jobs that they are contemplating creating. Our main point, relative to Pissarides (2009) and other standard labor search models, is to separate job-level and average variation in the wage for new hires, to map the model to our new evidence.

Our argument for the importance of job level wages does not hinge on the specific way in which we have modelled jobs. In Appendix Section E, we consider an alternative version of the Diamond-Mortensen-Pissarides model in which job-level wages are equally important. This alternative model features a continuum of firms with heterogeneous productivity and decreasing returns to scale in production, similar to Michaillat (2012), Acemoglu and Hawkins (2014) and Elsby and Michaels (2013), albeit with several simplifications relative to the latter two models.

Our baseline model has random search, but job-level wages are equally important in models with posted wages and directed search within job types. If there is directed search within job types, a similar equation to (16) will hold. This similarity is well known (e.g. Moen, 1997; Chodorow-Reich and Karabarbounis, 2016). In practice, US labor markets feature both posted and bargained wages, and both random and directed search (Hall and Krueger, 2012).

6.2.1 Unemployment is More Sensitive to Contractions in Labor Demand

Given that job-level wage changes are allocative for unemployment fluctuations, downward wage rigidity at the job level leads to asymmetry—unemployment is more sensitive to contractions than expansions in labor demand. We plug wage setting equations (8) and (9) into equation (16) to arrive at this result.

Corollary. *Starting from the steady state at time $t-1$, changes in unemployment satisfy to a first order*

$$\Delta \log u_t = \begin{cases} (-A + \overbrace{B\gamma}^{\substack{\text{wage}\\ \text{flexibility}\\ \text{upwards}}})\Delta \log y_t & \text{if } \Delta \log y_t > 0 \\ -A \Delta \log y_t & \text{if } \Delta \log y_t < 0 \end{cases} \tag{17}$$

Equation (17) makes a simple point. Unemployment is more sensitive to contractions than expansions in labor demand. During an expansion, so $\Delta \log y_t > 0$, the comovement between unemployment u_t and labor demand y_t is of lower magnitude. The differential sensitivity is entirely because the job-level wage for new hires is more rigid downwards. B captures the sensitivity of unemployment to wage changes. γ measures the extent to which wages move upwards with labor demand.

The intuition for this result is simple. After a contraction in output per worker, job-level wages do not fall. Firms' profits on new jobs fall sharply. Firms lower job creation, and unemployment rises. By contrast, after an expansion in output per worker, job-level wages rise. Firm's profits on new jobs increase more gradually—as a result, job creation and unemployment changes are smaller.

6.3 Calibration Exercise

We now do a simple calibration exercise. The goal of this exercise is to gauge the quantitative importance of our estimates of downward wage rigidity, for unemployment fluctuations.

We calibrate the parameters in equation (17). We then ask how much $\Delta \log u_t / \Delta \log y_t$ differs when the economy receives positive versus a negative labor demand shock. This difference summarizes how much downward wage rigidity matters for unemployment dynamics.

Table 11 contains the calibration values. Our calibration proceeds in two main steps.

In the first step, we calibrate the degree of wage flexibility upward, γ, using our new job level estimates. In section 4, we estimated the semielasticity of the wage for new hires with respect to unemployment, $\Delta \log w / \Delta U$, during expansions. To uncover $\Delta \log w / \Delta \log y$, we rescale $\Delta \log w / \Delta U$ by an estimate of $\Delta U / \Delta \log y$. The empirical counterpart to y is labor productivity, which we measure at the state-quarter level from the BEA's regional economic accounts. We estimate $\Delta U / \Delta \log y$ at quarterly frequency in an auxiliary regression, which we report in Appendix Section B, Table 16.[50] . The value of $\Delta \log w / \Delta \log y$ pins down wage flexibility upward, because γ is the elasticity of wages to a positive labor demand shock.[51] Second, we calibrate A and B using standard values from the labor search literature. For simplicity, we assume that $\bar{u}_H \approx \bar{u}_L, \phi_H \approx \phi_L$, to avoid taking a stand on the precise definition of a job type.

Our calibration reveals two results. First, downward wage rigidity leads to substantial asymmetry in unemployment dynamics. Table 11 also records this result. Unemployment is almost four times more sensitive to a contraction in labor demand than an expansion in labor demand. Starting from steady state, after an increase in y of 1 percent, unemployment decreases by 0.8 percent. After a decrease in y of 1 percent, unemployment rises by 3 percent. This asymmetry is entirely because wages are rigid downward and flexible upward. So, the degree of downward wage rigidity in the data is quantitatively important.

[50] We did not detect asymmetry in the comovement between U and y.

[51] We measure both labor productivity and wages in nominal terms, given that regional prices are measured with substantial error.

This first result from the calibration is key to the contribution of our paper. Many papers have previously argued that downward wage rigidity for new hires should lead to nonlinearity and asymmetry in unemployment dynamics (Dupraz, Nakamura, and Steinsson, 2016; Chodorow-Reich and Wieland, 2017). But there is limited previous evidence on whether wage rigidity for new hires can generate quantitatively relevant asymmetries. Our calibration exercise shows that downward wage rigidity is, indeed, quantitatively important.

The second calibration result is that unemployment fluctuations that are on average roughly as large as the data. Table 11 also records this result. The average value of $\Delta \log u_t / \Delta \log y_t$ predicted by the model is -1.92. For comparison, the value of $\Delta \log u_t / \Delta \log y_t$ in US time series data for 1950-2018 is -1.9.[52] So, wage rigidity can resolve the "unemployment volatility puzzle" of Shimer (2005).[53]

This second result from the calibration is also key to the contribution of our paper. Many previous papers, such as Hall (2005a) and Gertler and Trigari (2009), argue that wage rigidity can rationalize unemployment fluctuations. But consensus on wage rigidity for new hires is elusive. So, direct evidence that wages for new hires are rigid enough to rationalize the unemployment fluctuations in the data is important.

7 Conclusion

We use a new dataset to show downward rigidity in the wage for new hires at the job level. We have three findings. First, the wage for new hires 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 imply a downward constraint on the wage in newly created jobs. Second, at the job level, the wage for new hires rises during expansions but does not fall during contractions. So, wages are rigid downward and flexible upward at the job level. Third, wage flexibility upward is state dependent, consistent with downward rigidity.

One important question that our paper does not answer is *why* the wage for new hires is more rigid downward than upward at the job level. Several plausible mechanisms for downward wage rigidity largely apply to continuing workers and not for new hires. For example, firms might offer implicit contracts in the form of downwardly rigid wages to continuing workers, and not extend the same insurance to new hires (Beaudry and DiNardo, 1991). In ongoing work, we

[52]We calculate $\Delta \log u_t / \Delta \log y_t$ from a regression of the growth in quarterly unemployment on the growth in quarterly real output per worker, using the measures of quarterly unemployment and quarterly real output per worker from Table 11, without detrending either measure.

[53]Note that we calibrate ϕ to target the labor share. So, our model implies a substantially larger profit share than Hagedorn and Manovskii (2008). Hazell (2019) argues that the relevant notion of the profit share in DMP models includes the return to capital, consistent with our calibration.

seek to understand the mechanisms behind downward rigidity for new hires.

References

ABADIE, A., S. ATHEY, G. W. IMBENS, AND J. WOOLDRIDGE (2017): “When Should You Adjust Standard Errors for Clustering?,” Discussion paper, National Bureau of Economic Research.

ACEMOGLU, D., A. FINKELSTEIN, AND M. J. NOTOWIDIGDO (2013): “Income and health spending: Evidence from oil price shocks,” *Review of Economics and Statistics*, 95(4), 1079–1095.

ACEMOGLU, D., AND W. B. HAWKINS (2014): “Search with multi-worker firms,” *Theoretical Economics*, 9(3), 583–628.

ACHARYA, S., J. BENGUI, K. DOGRA, AND S. L. WEE (2018): “Slow recoveries and unemployment traps: monetary policy in a time of hysteresis,” .

ADAO, R., M. KOLESÁR, AND E. MORALES (2018): “Shift-share designs: Theory and inference,” Discussion paper, National Bureau of Economic Research.

ALLCOTT, H., AND D. KENISTON (2017): “Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America,” *The Review of Economic Studies*, 85(2), 695–731.

ANGRIST, J. D., AND A. B. KRUEGER (1995): “Split-sample instrumental variables estimates of the return to schooling,” *Journal of Business & Economic Statistics*, 13(2), 225–235.

BARATTIERI, A., S. BASU, AND P. GOTTSCHALK (2014): “Some evidence on the importance of sticky wages,” *American Economic Journal: Macroeconomics*, 6(1), 70–101.

BARLEVY, G. (2002): “The sullying effect of recessions,” *The Review of Economic Studies*, 69(1), 65–96.

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

BASU, S., AND C. L. HOUSE (2016): “Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models,” in *Handbook of Macroeconomics*, vol. 2, pp. 297–354. Elsevier.

BEAUDRY, P., AND J. 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), 665–688.

BERAJA, M., E. HURST, AND J. OSPINA (2016): “The aggregate implications of regional business cycles,” Discussion paper, National Bureau of Economic Research.

BEWLEY, T. F. (2002): *Why wages don't fall during a recession*. Harvard University Press.

BILS, M., Y. CHANG, AND S.-B. KIM (2014): "How sticky wages in existing jobs can affect hiring," Discussion paper, National Bureau of Economic Research.

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

BOUND, J., AND A. B. KRUEGER (1991): "The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?," *Journal of Labor Economics*, 9(1), 1–24.

CACCIATORE, M., AND F. RAVENNA (2020): "Uncertainty, wages, and the business cycle," .

CAJNER, T., AND D. RATNER (2016): "A Cautionary Note on the Help Wanted Online Data," *FEDS Notes, Board of Governors of the Federal Reserve System https://www. federalreserve. gov/econresdata/notes/feds-notes/2016/acautionary-note-on-the-help-wanted-online-data-20160623. html*.

CARD, D., AND D. HYSLOP (1997): "Does inflation" grease the wheels of the labor market"?," in *Reducing inflation: Motivation and strategy*, pp. 71–122. University of Chicago Press.

CARNEIRO, A., P. GUIMARÃES, AND P. 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, A. P., T. JAYASUNDERA, AND D. REPNIKOV (2014): "Understanding online job ads data," *Georgetown University, Center on Education and the Workforce, Technical Report (April)*.

CARRILLO-TUDELA, C., H. GARTNER, AND L. KAAS (2020): "Recruitment Policies, Job-Filling Rates and Matching Efficiency," .

CHODOROW-REICH, G., AND L. KARABARBOUNIS (2016): "The cyclicality of the opportunity cost of employment," *Journal of Political Economy*, 124(6), 1563–1618.

CHODOROW-REICH, G., AND J. WIELAND (2017): "Secular labor reallocation and business cycles," Discussion paper, National Bureau of Economic Research.

CHOI, S., N. FIGUEROA, AND B. VILLENA-ROLDÁN (2020): "Wage Cyclicality Revisited: The Role of Hiring Standards," .

DALY, M. C., AND B. HOBIJN (2014): "Downward nominal wage rigidities bend the Phillips curve," *Journal of Money, Credit and Banking*, 46(S2), 51–93.

DAVIS, S. J., R. J. FABERMAN, AND J. C. HALTIWANGER (2013): “The establishment-level behavior of vacancies and hiring,” *The Quarterly Journal of Economics*, 128(2), 581–622.

DONIGER, C. (2019): “Do Greasy Wheels Curb Inequality?,” .

DUPRAZ, S., E. NAKAMURA, AND J. STEINSSON (2016): “A plucking model of business cycles,” Discussion paper, mimeo.

ELIAZ, K., AND R. SPIEGLER (2014): “Reference dependence and labor market fluctuations,” *NBER macroeconomics annual*, 28(1), 159–200.

ELSBY, M. W. (2009): “Evaluating the economic significance of downward nominal wage rigidity,” *Journal of Monetary Economics*, 56(2), 154–169.

ELSBY, M. W., AND R. MICHAELS (2013): “Marginal jobs, heterogeneous firms, and unemployment flows,” *American Economic Journal: Macroeconomics*, 5(1), 1–48.

ELSBY, M. W., D. SHIN, AND G. SOLON (2016): “Wage adjustment in the Great Recession and other downturns: Evidence from the United States and Great Britain,” *Journal of Labor Economics*, 34(S1), S249–S291.

ELSBY, M. W., AND G. SOLON (2018): “How Prevalent Is Downward Rigidity in Nominal Wages? International Evidence from Payroll Records and Pay Slips,” Discussion paper, National Bureau of Economic Research.

GAVAZZA, A., S. MONGEY, AND G. L. VIOLANTE (2018): “Aggregate recruiting intensity,” *American Economic Review*, 108(8), 2088–2127.

GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2016): “Unemployment fluctuations, match quality, and the wage cyclicality of new hires,” Discussion paper, National Bureau of Economic Research.

GERTLER, M., AND A. TRIGARI (2009): “Unemployment fluctuations with staggered Nash wage bargaining,” *Journal of political Economy*, 117(1), 38–86.

GRIGSBY, J. (2019): “Skill Heterogeneity and Aggregate Labor Market Dynamics,” Discussion paper, Working Paper, University of Chicago 2019.

GRIGSBY, J., E. HURST, AND A. YILDIRMAZ (2018): “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data,” *Working paper*.

HAEFKE, C., M. SONNTAG, AND T. VAN RENS (2013): “Wage rigidity and job creation,” *Journal of monetary economics*, 60(8), 887–899.

HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2013): "Unemployment benefits and unemployment in the great recession: the role of macro effects," Discussion paper, National Bureau of Economic Research.

HAGEDORN, M., AND I. MANOVSKII (2008): "The cyclical behavior of equilibrium unemployment and vacancies revisited," *American Economic Review*, 98(4), 1692–1706.

——— (2013): "Job selection and wages over the business cycle," *American Economic Review*, 103(2), 771–803.

HALL, R. E. (2005a): "Employment fluctuations with equilibrium wage stickiness," *American economic review*, 95(1), 50–65.

——— (2005b): "Job loss, job finding, and unemployment in the US economy over the past fifty years," *NBER macroeconomics annual*, 20, 101–137.

HALL, R. E., AND A. 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, R. E., AND P. R. MILGROM (2008): "The limited influence of unemployment on the wage bargain," *American economic review*, 98(4), 1653–74.

HARRIS, M., AND B. HOLMSTROM (1982): "A theory of wage dynamics," *The Review of Economic Studies*, 49(3), 315–333.

HAZELL, J. (2019): "The Falling Labor Share Dampens Unemployment Fluctuations," *Mimeo*.

HECKMAN, J., AND B. SINGER (1984): "A method for minimizing the impact of distributional assumptions in econometric models for duration data," *Econometrica: Journal of the Econometric Society*, pp. 271–320.

HECKMAN, J. J. (1979): "Sample selection bias as a specification error," *Econometrica: Journal of the econometric society*, pp. 153–161.

HERSHBEIN, B., AND L. B. KAHN (2016): "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings," Discussion paper, National Bureau of Economic Research.

JÄGER, S., B. SCHOEFER, S. G. YOUNG, AND J. ZWEIMÜLLER (2018): "Wages and the Value of Nonemployment," Discussion paper, National Bureau of Economic Research.

JARDIM, E. S., G. SOLON, AND J. L. VIGDOR (2019): "How Prevalent Is Downward Rigidity in Nominal Wages? Evidence from Payroll Records in Washington State," Discussion paper, National Bureau of Economic Research.

JO, Y. J. (2019): "Downward nominal wage rigidity in the United States," Discussion paper, Technical report, Columbia University.

KAUR, S. (2019): "Nominal wage rigidity in village labor markets," *American Economic Review*, 109(10), 3585–3616.

KEYNES, J. M. (1936): *The general theory of employment, interest, and money*. Springer.

KILIAN, L., AND R. J. VIGFUSSON (2011): "Are the responses of the US economy asymmetric in energy price increases and decreases?," *Quantitative Economics*, 2(3), 419–453.

KLENOW, P. J., AND O. KRYVTSOV (2008): "State-dependent or time-dependent pricing: Does it matter for recent US inflation?," *The Quarterly Journal of Economics*, 123(3), 863–904.

KUDLYAK, M. (2014): "The cyclicality of the user cost of labor," *Journal of Monetary Economics*, 68, 53–67.

KURMANN, A., AND E. MCENTARFER (2017): "Downward Wage Rigidity in the United States: New Evidence from Administrative Data," .

LJUNGQVIST, L., AND T. J. SARGENT (2017): "The fundamental surplus," *American Economic Review*, 107(9), 2630–65.

MAKRIDIS, C., AND M. GITTLEMAN (2019): "On the Cyclicality of Real Wages and Employment: New Evidence and Stylized Facts from Performance Pay and Fixed Wage Jobs," *Available at SSRN 3017034*.

MARINESCU, I., AND R. WOLTHOFF (2016): "Opening the black box of the matching function: The power of words," Discussion paper, National Bureau of Economic Research.

MARTINS, P. S., G. SOLON, AND J. 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.

MENZIO, G., AND S. SHI (2011): "Efficient search on the job and the business cycle," *Journal of Political Economy*, 119(3), 468–510.

MICHAILLAT, P. (2012): "Do matching frictions explain unemployment? Not in bad times," *American Economic Review*, 102(4), 1721–50.

MOEN, E. R. (1997): "Competitive search equilibrium," *Journal of political Economy*, 105(2), 385–411.

MORTENSEN, D. T., AND C. A. PISSARIDES (1994): "Job creation and job destruction in the theory of unemployment," *The review of economic studies*, 61(3), 397–415.

MUELLER, A. I. (2017): "Separations, Sorting, and Cyclical Unemployment," *American Economic Review*, 107(7), 2081–2107.

NAKAMURA, E., AND J. STEINSSON (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

NAKAMURA, E., AND J. STEINSSON (2014): "Fiscal stimulus in a monetary union: Evidence from US regions," *American Economic Review*, 104(3), 753–92.

OF ATLANTA, F. R. B. (2014): "Wage Growth A Missing Piece of the Full Employment Puzzle," *https://www.frbatlanta.org/about/publications/annual-reports/2014/wage-growth-missing.aspx*.

PETRONGOLO, B., AND C. A. PISSARIDES (2001): "Looking into the black box: A survey of the matching function," *Journal of Economic literature*, 39(2), 390–431.

PETROSKY-NADEAU, N., AND L. ZHANG (2013): "Solving the DMP model accurately," Discussion paper, National Bureau of Economic Research.

PETROSKY-NADEAU, N., L. ZHANG, AND L.-A. KUEHN (2018): "Endogenous disasters," *American Economic Review*, 108(8), 2212–45.

PISSARIDES, C. A. (2009): "The unemployment volatility puzzle: Is wage stickiness the answer?," *Econometrica*, 77(5), 1339–1369.

ŞAHIN, A., J. SONG, G. TOPA, AND G. L. VIOLANTE (2014): "Mismatch unemployment," *American Economic Review*, 104(11), 3529–64.

SCHAEFER, D., AND C. SINGLETON (2019): "Cyclical labor costs within jobs," *European Economic Review*, 120, 103317.

SCHMITT-GROHÉ, S., AND M. URIBE (2016): "Downward nominal wage rigidity, currency pegs, and involuntary unemployment," *Journal of Political Economy*, 124(5), 1466–1514.

SCHOEFER, B. (2015): "The financial channel of wage rigidity," Ph.D. thesis, Harvard University.

SHIMER, R. (2005): "The cyclical behavior of equilibrium unemployment and vacancies," *American economic review*, 95(1), 25–49.

SOLON, G., R. BARSKY, AND J. A. PARKER (1994): "Measuring the cyclicality of real wages: how important is composition bias?," *The quarterly journal of economics*, 109(1), 1–25.

YAGAN, D. (2016): "Is the great recession really over? longitudinal evidence of enduring employment impacts," *Unpublished Manuscript*.

8 Tables

Table 1: Summary Statistics

	Min	Max	Average	
Posts Per State	4799	3012689	421412	
Posts Per Quarter	279252	1278327	782622	
Posts Per State-Quarter	49	190582	15050	
Posts Per 6 Digit SOC Code	1	1925439	25500	
Total Posts	21913422			
Share Missing Job Title	.57			
Share Missing Establishment Code	.57			
Share of 6 digit SOC occupations covered	.99			
Share posting wage range	.44			
Average width of range	.077			
Pay Categories:				
	Base Pay	Bonus	Total Pay	**Total**
Annual	3962172	530169	3648138	8234372
Daily	330899	306405	857674	1520389
Hourly	6067618	376666	3918815	10397359
Monthly	380414	438023	743509	1577650
Weekly	80038	22368	68401	183652
Total	10821141	1673631	9236537	21913422

Notes: the width of the wage range is defined as (Max – Min) / Max, where Max and Min are the maximum and minimum of the wage range. The share of 6 digit SOC occupations covered, is defined as the share of 6 digit occupations in the 2014-2016 Occupational Employment Statistics (OES) that post in Burning Glass, weighted by OES employment.

Table 2: Comparison of Wage for New Hires in CPS and BG, by State-Quarter

Panel A:	Log CPS New Hire Wage, by State-Quarter					
	(1)	(2)	(3)	(4)	(5)	(6)
Independent Variable:						
Log Burning Glass Wage,	0.970***	1.034***	0.715***			
All Vacancies	(0.174)	(0.252)	(0.108)			
Log Burning Glass Wage,				0.957***	1.017***	0.706***
Vacancies with Job Code only				(0.171)	(0.246)	(0.106)

Panel B:	Log QWI New Hire Earnings, by State-Quarter					
	(1)	(2)	(3)	(4)	(5)	(6)
Independent Variable:						
Log Burning Glass Wage,	1.246***	1.184***	1.007***			
All Vacancies	(0.203)	(0.347)	(0.140)			
Log Burning Glass Wage,				1.234***	1.168***	0.997***
Vacancies with Job Code only				(0.201)	(0.341)	(0.139)
State Effects	N	Y	N	N	Y	N
Time Effects	N	N	Y	N	N	Y
Number of Observations	1428	1428	1428	1428	1428	1428
State Clusters	52	52	52	52	52	52

Notes: in Panel A, the dependent variable is the log of the hours-weighted mean wage for newly hired workers from the 2010-2016 CPS, by quarter and state. Newly hired workers are identified using the rotating panel structure of the Basic Monthly File, and wages are from the Outgoing Rotation Group. Wages are trimmed at the first and 99th percentile. The wage is usual hourly earnings for hourly and non-hourly workers, constructed following CEPR's "wage 3" series, for non-farm workers. The regression in panel A is weighted by the number of CPS observations in each state-quarter.

In Panel B, the dependent variable is the log of the mean hourly earnings for newly hired workers from the 2010-2016 Quarterly Workforce Indicators (QWI), by quarter and state. The regression in panel B is weighted by the number of hires in the quarter, also from the QWI.

In the 2010-2016 Burning Glass, the wage is the log of workers' salaries. Salaries are reported by pay frequency (e.g. hourly or annual) and salary type (e.g. base pay or total pay). Salaries are trimmed at the 5th and 95th percentiles in each year, within each pay frequency and salary type. To uncover state-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 state by quarter dummies. Observations are weighted by the 2014-2016 OES. Then W_{st} is the log mean salary in the state-quarter, after adjusting for pay frequency and salary type. We split the sample in half in each state-quarter, and instrument for salaries in one sub-sample with salaries in the other, to overcome measurement error. A vacancy has a job code if it has a non-missing establishment and job title identifier.

Standard errors are two way clustered by quarter and state. One, two and three asterisks denote significance at the 5, 1 and 0.1 percent levels, respectively.

Table 3: Summary Statistics, for Data Differenced by Job

	Min	Max	Average	Total
Total Vacancy Posts				1598505
Share of 6 digit SOC occupations covered in the OES				.99
Posts Per Job	2	23	2.5	
Jobs per 6 digit SOC occupation	1	176081	1247.2	
Jobs per State	264	118076	19909	
Jobs per Quarter	7519	117566	38343	

Notes: a job is an employer by location by pay frequency by salary type by job title unit. We take the quarterly average wage by job, and then difference by the job.

Table 4: Quarterly Job-Level Statistics On Wage for New Hires

	Unweighted (1)	OES Weights (2)	QCEW Weights (3)	High Wage Jobs (4)
Probability of Job-Level Wage Change	0.17	0.16	0.16	0.16
Probability of Job-Level Wage Increase	0.11	0.11	0.11	0.11
Probability of Job-Level Wage Decrease	0.04	0.04	0.04	0.04
Implied Duration for which Job-Level Wages Are Unchanged (Quarters)	5.45	5.77	5.53	5.46
N	1598505	1598505	1598505	1198879

Notes: a job is an establishment by region by job title by salary type by pay frequency observation. The wage for new hires is averaged within each job-quarter. The sample is the 2010-2016 Burning Glass data. We estimate the probability of job-level wage change using a similar method to Klenow and Kryvtsov (2008) and 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 column (2), we reweight to target the distribution of jobs at the 6 digit SOC level from the 2014-2016 OES. In column (3) we reweight to target the distribution of employment across states from the 2010 QCEW. In column (4) we drop jobs in the bottom quartile of the wage distribution.

Table 5: Regression of Job-Level Wage Growth for New Hires on Unemployment Changes

Dependent Variable:	Quarterly Job-Level Growth in Wage for New Hires				
	(1)	(2)	(3)	(4)	(5)
Independent Variable:					
ΔU_{st}	-0.0517	0.152	0.0454	-0.946***	0.367
	(0.256)	(0.308)	(0.249)	(0.102)	(0.892)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.255***	-1.413***	-1.309***		-2.397*
	(0.265)	(0.371)	(0.246)		(1.029)
Job-Level Difference	Y	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y	Y
State Trend	N	Y	N	N	N
OES Weight	N	N	Y	N	N
Oil Shock IV	N	N	N	N	Y
N	1566182	1566182	1511642	1566182	1566182
State Clusters	52	52	52	52	52

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW, in all but the final column. In columns (1)-(3), we project positive and negative unemployment changes on positive and negative employment growth changes. We control for an indicator variable for whether unemployment changes are positive or negative.

In the final column, we instrument for unemployment with a Bartik-style instrument based on the oil price. The first stage regression is $\Delta U_{st} = \sum_s \left[\beta_s \Delta \log\left(\text{oil price}_{t-1}\right) + \gamma_s I\left(\Delta \log\left(\text{oil price}_{t-1}\right) < 0\right) \Delta \log\left(\text{oil price}_{t-1}\right)\right] + \text{error}_{st}$, where α_s, β_s and γ_s are estimated, similarly to Nakamura and Steinsson (2014). oil price_t is the price of Brent crude oil averaged over quarter t.

Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state; except for the last column, which clusters by state and quarter. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. In some specifications, we reweight to target the occupation employment distribution at the 6 digit SOC level from the 2014-2016 OES.

Table 6: Estimates of Downward Wage Rigidity—Robustness

Specification	Coefficient on $\Delta U_{st} \times I(\Delta U_{st} < 0)$	S.E.	N
1. Baseline	-1.255***	(0.265)	1566182
2. Excludes vacancies in consecutive quarters	-2.061***	(0.308)	626034
3. Annual	-2.703**	(0.959)	656596
4. Controls for time since last vacancy	-0.520**	(0.163)	1566182
5. Heckman correction for selection bias	-1.402***	(0.244)	1566182
6. QCEW Weighted	-1.161***	(0.257)	1566182
7. X11 adjustment	-4.366***	(1.536)	1566182
8. State ×Quarter-of-year FEs	-0.931*	(0.435)	1566182
9. No wage ranges	-1.174***	(0.303)	795316
10. No bonuses	-1.290***	(0.287)	1410347
11. Alternative job definition	-1.744***	(0.233)	1229020
12. No time fixed effect	-1.319***	(0.165)	1566182

Notes: the first row reports the coefficient on $\Delta U_{st} \times I(\Delta U_{st} < 0)$ from the benchmark regression, that is, column (1) of Table 5. The second row reports the coefficient from the benchmark regression, after excluding vacancies posted in the quarter immediately after another vacancy of the same job. The third row runs the baseline regression at annual frequency, by taking the mean wage of vacancies within a job and year. The fourth row controls for the time that has elapsed since the job last posted a vacancy. The fifth row corrects for selection bias following Heckman (1979). The selection equation includes the regressors in the baseline regression equation as well as the level of quarterly state log employment and unemployment equation. The sixth row reweights to target mean employment in each state over 2010-2016, from the Quarterly Census of Employment and Wages. The seventh row reports the coefficient from the benchmark regression, after seasonally adjusting using the Census Bureau's X-11 algorithm. We seasonally adjust at the state-quarter level for 1980-2017 data, for both unemployment and employment. The eighth row reports the coefficient from the benchmark regression, after also controlling for the interaction of quarter-of-year and state fixed effects. The ninth row drops from the sample all vacancies that post a range of wages, instead of a point wage. The tenth row excludes vacancies with bonus pay. The eleventh row uses an alternative definition of a job, by taking the mean wage across job titles by establishments, averaging over workers paid at different frequencies (e.g. averaging over hourly and annual paid workers). A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. Standard errors are clustered by state.

Table 7: Regression of Establishment Wages for New Hires on Unemployment

Dependent Variable:	Quarterly Establishment-Level Growth in Wage for New Hires			
	(1)	(2)	(3)	(4)
Independent Variable:				
ΔU_{st}	0.00392	-0.268	-0.0431	-0.909***
	(0.313)	(0.353)	(0.341)	(0.0737)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.082**	-0.785^{+}	-1.021*	
	(0.382)	(0.427)	(0.414)	
$(\Delta U_{st})^2$				
Establishment-Level Difference	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y
State Trend	N	Y	N	N
QCEW Weight	N	N	Y	N
N	1845695	1845695	1845695	1845695
State Clusters	52	52	52	52

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each establishment-quarter, separately for each pay frequency (e.g. hourly, monthly or annual) and salary type (e.g. base pay or total pay). The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. In the asymmetric specifications, we project positive and negative unemployment changes on positive and negative employment growth changes and control for an indicator variable for whether unemployment changes are positive or negative. Wage growth is trimmed at the 1st and 99th percentiles. An establishment is an employer by location by pay frequency by salary type unit. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. In some specifications, we reweight to target the average regional employment distribution at the state level from the 2010-2016 QCEW.

Table 8: Wage Rigidity After Contractions and Expansions

Dependent Variable:	Quarterly Job-Level Growth in Wage for New Hires			
	(1)	(2)	(3)	(4)
Independent Variables:				
ΔU_{st}	-0.557**	-0.486**	-0.621***	-0.559**
	(0.184)	(0.168)	(0.0555)	(0.185)
$\Delta U_{st} \times$	-0.727***	-0.744***	-0.646***	-0.725***
$I\left(U_{s,t-1} - U_{s,t-13} < 0\right)$	(0.138)	(0.130)	(0.0561)	(0.139)
Job-Level Difference	Y	Y	Y	Y
Time Effects	Y	Y	Y	Y
State Trend	N	Y	N	N
State $\times$Quarter-of-Year Effects	N	N	Y	N
Level of Unemployment	N	N	N	Y
Number of Observations	1089785	1089785	1089785	1089785
State Clusters	52	52	52	52

Notes: We estimate the regression

$$\Delta \log w_{jst} = \alpha + \gamma_t + \kappa \Delta U_{st} + \nu \Delta U_{st} \times I\left(U_{s,t-1} - U_{s,t-13} < 0\right) + \varepsilon_{jst}$$

The dependent variable is quarterly job-level wage growth, in percentage points, for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, also in percentage points. We interact state-quarter unemployment changes with an indicator for whether state unemployment fell over the previous three years, and we also add a control for this indicator variable. We restrict the sample only to observations for which $\Delta U_{st} < 0$. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW, employment is interacted with the same indicator. Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 9: Regression of State Share in High Wage Vacancies on Unemployment

Panel A:	Quarterly Change in State Share of High Wage Vacancies			
	(1)	(2)	(3)	(4)
ΔU_{st}	-0.654	-1.040	4.815	-0.0414
	(0.831)	(1.286)	(2.677)	(0.393)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	0.982	1.549	-3.537	
	(1.270)	(1.927)	(5.138)	
State Difference	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y
State Trend	N	Y	N	N
QCEW Weight	N	N	Y	N
N	1404	1404	1404	1404
State Clusters	51	51	51	51

Notes: the dependent variable is the change in the quarterly share of high wage vacancies within each state. High wage vacancies have a wage above the national median wage, by salary type and pay frequency, in 2010-2016 Burning Glass. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. In the asymmetric specifications, we project positive and negative unemployment changes on positive and negative employment growth changes. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia. In some specifications, we reweight to target the average regional employment distribution at the state level from the 2010-2016 QCEW, otherwise we weight by state number of vacancies in Burning Glass.

Table 10: Estimates of Downward Wage Rigidity from Average and Worker-Level Wages

	Quarterly Growth in Average Wage for New Hires					
	State Average				National Average	
	Burning Glass		CPS		CPS	NLSY
	(1)	(2)	(3)	(4)	(5)	(6)
ΔUnemployment	-0.0226	-0.275	-3.379	-3.888	3.770	-1.779
	(1.543)	(0.903)	(3.227)	(3.486)	(3.468)	(3.172)
ΔUnemployment ×	1.791	1.627	3.255	1.291	-5.108	2.935
$I\left(\Delta\text{Unemployment} < 0\right)$	(1.852)	(1.382)	(3.191)	(2.903)	(5.151)	(4.311)
Occupation Controls	N	Y	N	Y	N	N
Industry Controls	N	Y	Y	Y	N	Y
Worker Demographic Controls	N	N	Y	Y	Y	Y
Worker Fixed Effects	N	N	N	N	N	Y
Hagedorn and Manovskii (2013) Control	N	N	N	N	N	Y
N	1377	1377	1377	918	83	83

Notes: each column regresses a measure of wage growth for new hires on unemployment changes. In the first and second columns, the dependent variable is the percentage growth in average state wages, from Burning Glass. The independent variables are state unemployment changes, an indicator if state unemployment is falling, and the interaction of the indicator with state unemployment. We project the dependent variables onto state-quarter employment growth from the 2010-2016 QCEW, and interact employment growth with an indicator for whether employment growth is positive. The sample period is 2010-2016, the sample is vacancies in the 50 states plus the District of Columbia. In column (1), average state wages in Burning Glass are measured in the same way as in Table 2. Column (2) is the average of residual state wages, controlling for 2 digit industry fixed effects and 6 digit occupation fixed effects. In the third and fourth columns, the dependent variable is the percentage growth in state wages for newly hired workers, from the CPS. The independent variables and sample details are the same as in columns (1) and (2). Wages for new hires in the CPS are measured in the same way as in Table 2. Column (3) is the average of residual state wages, controlling for worker-specific observable characteristics, namely gender, race, marital status, education, a fourth order polynomial in experience and census industry code fixed effects in column (3). Column (4) additional controls for census occupation code fixed effects, consistent occupation codes are only available after 2012. In the fifth column, the dependent variable is the quarterly percentage growth in the national median wage for workers newly hired from unemployment. This wage series is for 1984-2006, and is taken from Haefke, Sonntag, and Van Rens (2013). The wage series is the average of the residual after regressing wages on worker-specific observable characteristics. We regress wage growth on the change in national unemployment, an indicator for whether national unemployment is falling, and the interaction of the change in national unemployment with the indicator. In the the sixth column, the dependent variable is the quarterly percentage growth in the national mean wage for newly hired workers. This wage series is for 1984-2006, and is measured from the National Longitudinal Survey of Youth, from Basu and House (2016). This specification also controls for job composition using cumulative tightness, as in Hagedorn and Manovskii (2013). The wage series controls for worker-specific fixed effects, and age. Standard errors are clustered by state in the first four columns, and are heteroskedasticity robust in the last two columns. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively.

Table 11: Calibration

Calibration Inputs

Parameter	Meaning	Value	Source
α	Matching function parameter	0.5	Petrongolo and Pissarides (2001), Şahin, Song, Topa, and Violante (2014)
u	Average unemployment	5.8%	2000-2018 average
τ	Quarterly separation rate	3.5%	2000-2018 average
ρ	Quarterly autocorrelation of output per worker	0.98	2000-2018 average
β	Quarterly discount rate	0.99	Shimer (2005)
ϕ	Labor share	0.78	2000-2018 labor share
γ	Wage flexibility upwards	0.65	Our estimates

Calibration Result

	Labor demand falling, $\Delta \log y < 0$	Labor demand rising, $\Delta \log y > 0$	Average Value
$\Delta \log u_t / \Delta \log y_t$	-3.00	-0.83	-1.92

Notes: the unemployment series is from the BLS's series, derived from the Current Population Survey. The separation rate is from JOLTS. The labor share is the ratio (compensation / (compensation + net value added)) for the non-financial corporate sector, as recorded in NIPA. The series for quarterly output per worker is real GDP, from the Bureau of Economic Analysis, divided by total non farm employment from the Bureau of Labor Statistics. We calculate the first autocorrelation after removing a quadratic trend in output per worker, estimated over 1950-2018.

Appendix

A Additional Figures

Figure 1: Binned Scatter of State-Quarter Wages for New Hires, in the CPS and Burning Glass

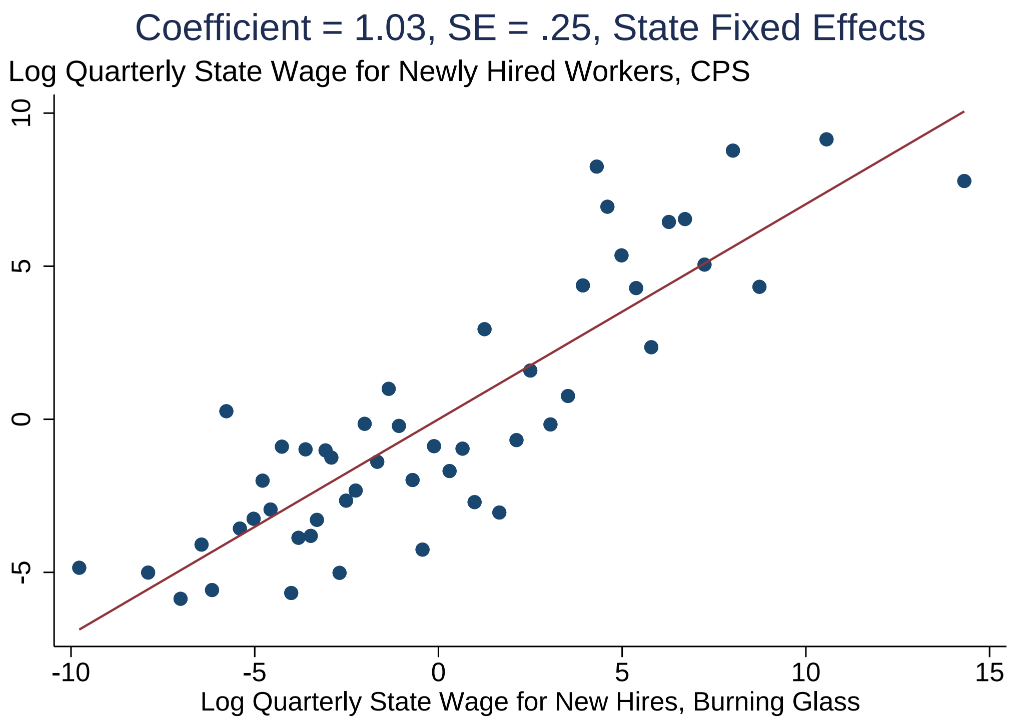

Notes: the y axis variable is the log of hours-weighted mean state-quarter wage for newly hired workers, from the 2010-2016 CPS. The x axis variable is the log of mean state-quarter wages for new hires, from the 2010-2016 Burning Glass data. The graph plots the weighted mean value of the y variable, for 50 equally sized weighted bins of the x variable. Bins and means are weighted by the size of each state-quarter in the CPS. The line is from a least squares regression, weighted the same way, the standard error is clustered by state. Mean state-quarter wages for new hires in the CPS, and for new hires in Burning Glass, are calculated in the same way as in table 2.

Figure 2: Comparison of Employment Shares by Occupation, in Burning Glass and the OES

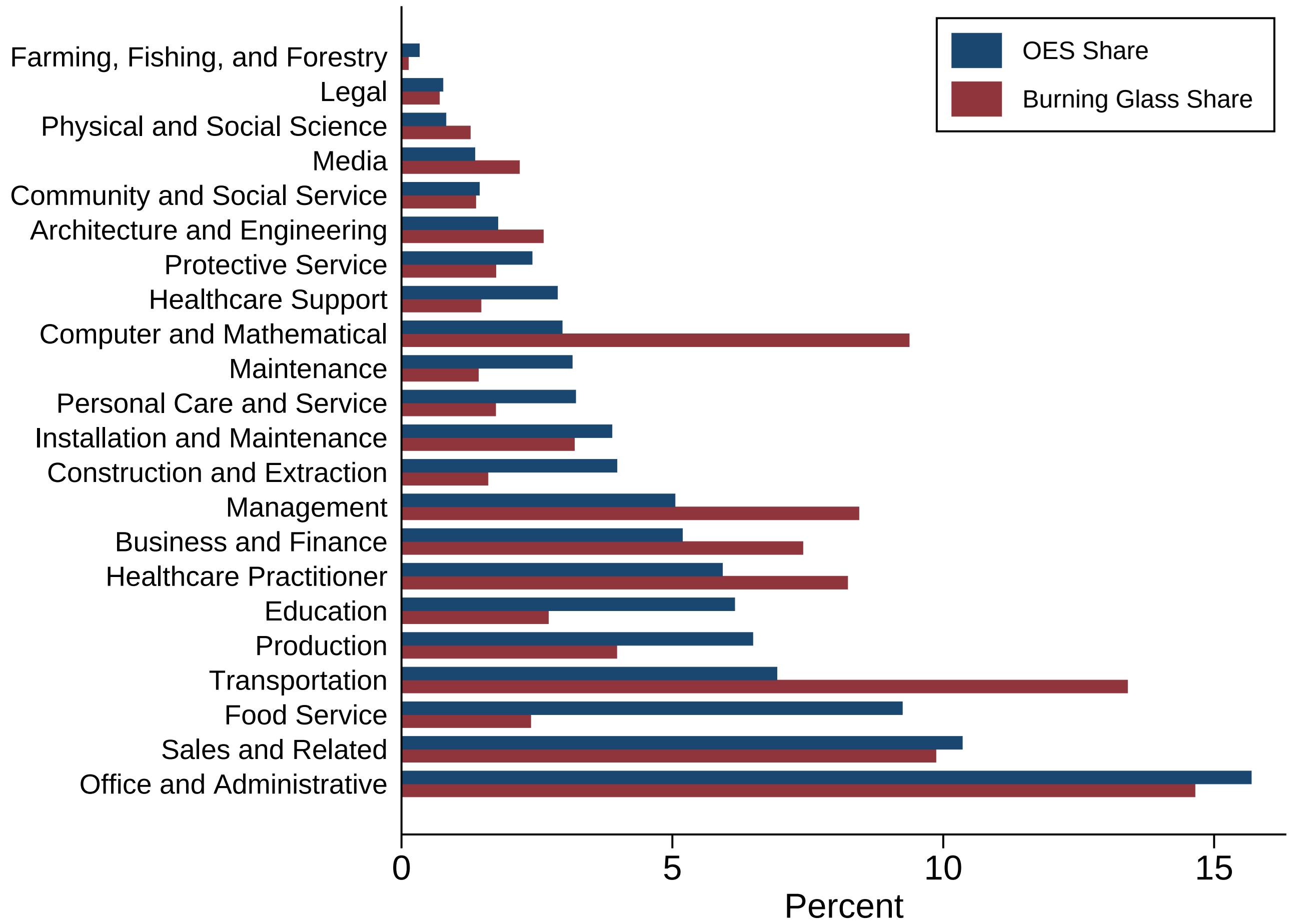

Notes: In Burning Glass, the data is 2010-2016; in the OES, the data is 2014-2016. In both datasets, the comparison is at the 2 digit SOC level, and excludes military.

Figure 3: Share of Wage Growth Observations in Each State with Falling Unemployment

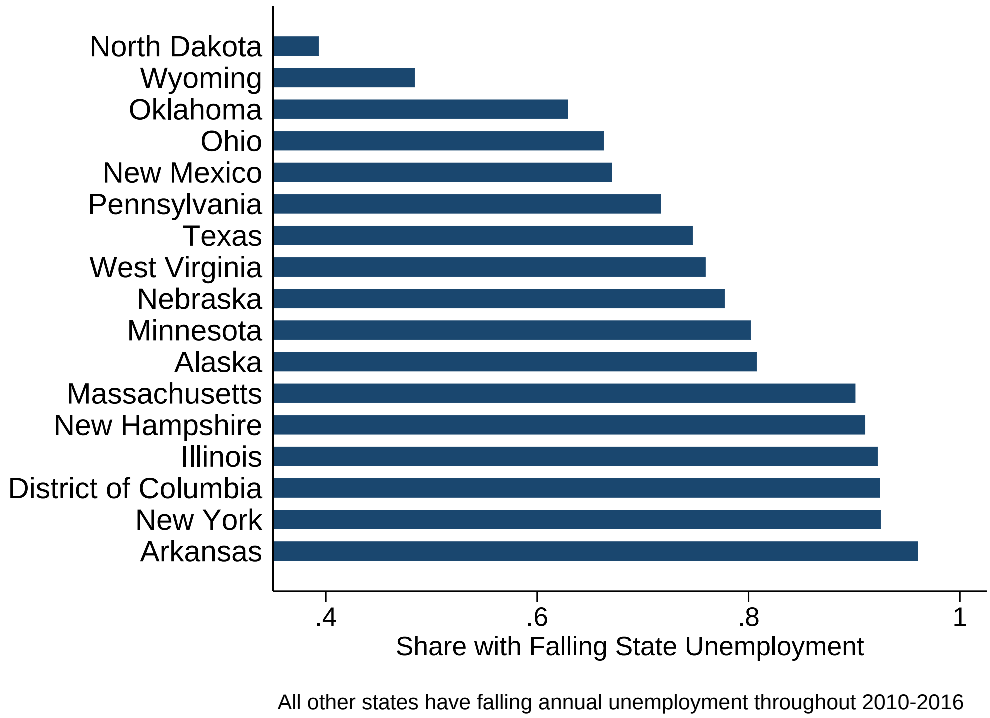

Notes: this graph plots the share of wage growth observations in each state, for which annual state unemployment is falling during the year of the wage posting, for the state in which the vacancy is posted. Log wages are differenced by job. The time period is 2010-2016. Unemployment is from the LAUS. Wages are averaged by job-year, where a job is a job title by establishment by salary type by pay frequency unit.

Figure 4: Share of Wage Growth Observations in Each Year with Falling Unemployment

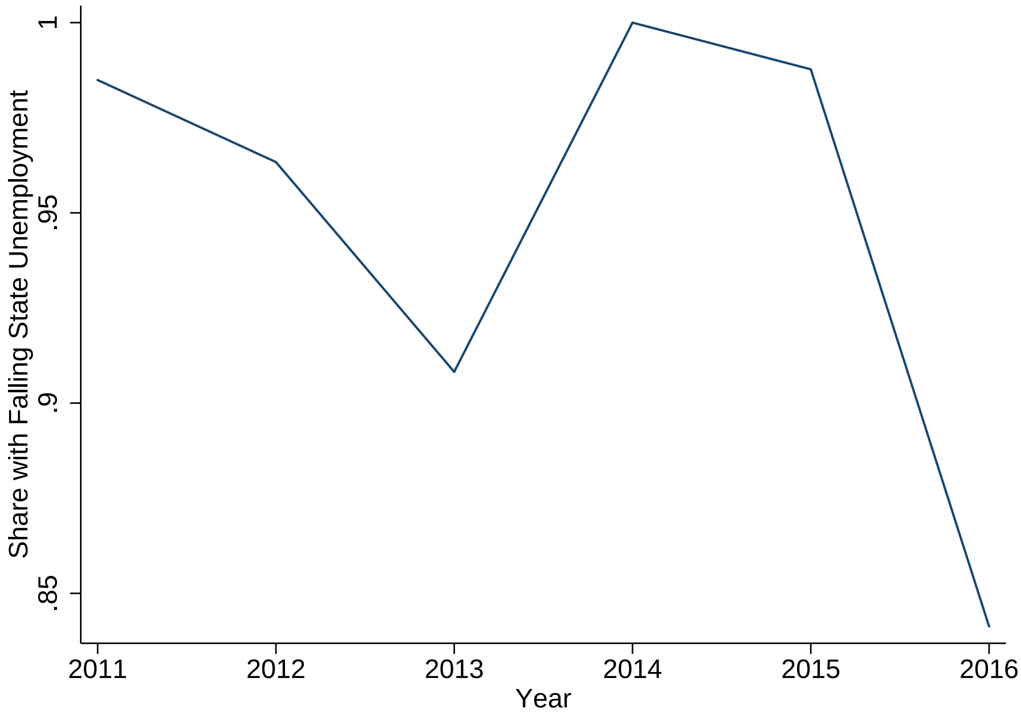

Notes: this graph plots the share of wage growth observations in each year, for which annual state unemployment is falling during the year of the wage posting, for the state in which the vacancy is posted. Log wages are differenced by job. The time period is 2010-2016. Unemployment is from the LAUS. Wages are averaged by job-year, where a job is a job title by establishment by salary type by pay frequency unit.

Figure 5: Quarterly Global Oil Prices

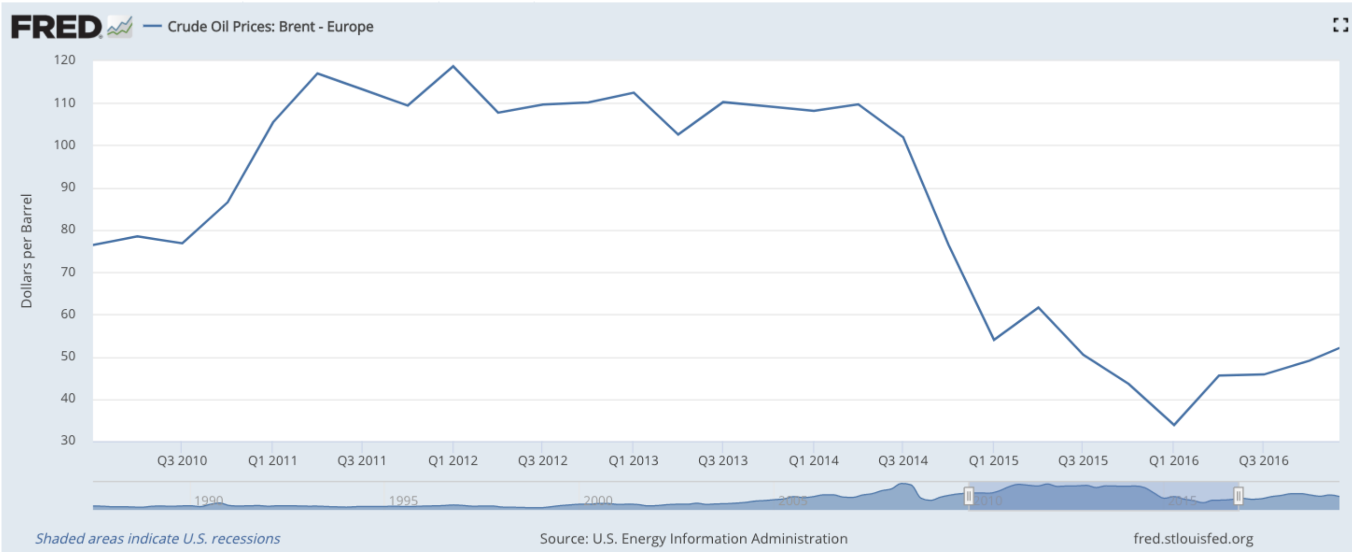

Notes: this figure plots the quarterly average oil price for 2010 to 2016, using the Brent Crude measure.

B Additional Tables

Table 1: Cyclicality of Whether Firms Include Wages In Vacancies

Dependent Variable:	Change in Share of State Vacancies with Wage			
	(1)	(2)	(3)	(4)
Quarterly State	0.0102	0.00746		
Unemployment Change	(0.0136)	(0.0214)		
Annual State			0.00304	-0.0111
Unemployment Change			(0.00542)	(0.0139)
State Difference	Y	Y	Y	Y
Time Effects	Y	Y	Y	Y
State Trend	N	Y	N	Y
Number of Observations	1377	1377	306	306
State Clusters	52	52	52	52

Notes: the dependent variable is the change in percentage points in the share of vacancies in the state that post a wage in the time period, from a 5% sample of the 2010-2016 Burning Glass dataset, inclusive of all vacancies that do or do not post wages. The independent variable is the change in percentage points in state-quarter or state-year unemployment from the 2010-2016 LAUS, projected onto employment growth from the 2010-2016 QCEW. Standard errors are in parentheses, clustered by state. One, two and three asterisks denote significance at the 5, 1 and 0.1 percent levels. Observations are weighed by 2010 state employment from the QCEW.

Table 2: Comparison of Burning Glass Wages with Occupation Wages and City Earnings

Dependent Variable:	Log Median Hourly Wage by Occupation (OES)			
	(1)	(2)	(3)	(4)
Independent Variable:				
Log Median Salary	1.139***	1.174***	0.779***	1.001***
by Occupation (BG)	(0.0945)	(0.0678)	(0.0883)	(0.0899)
BG Contract Type	Base Pay, Annual	Base Pay, Hourly	Total Pay, Annual	Total Pay, Hourly
Observations	742	751	742	754

Dependent Variable:	Log Average Weekly Earnings by CBSA (QCEW)			
	(1)	(2)	(3)	(4)
Independent Variable:				
Log Median Salary	1.295***	1.390***	1.069***	0.900***
by CBSA (BG)	(0.0754)	(0.127)	(0.100)	(0.149)
BG Contract Type	Base Pay, Annual	Base Pay, Hourly	Total Pay, Annual	Total Pay, Hourly
Observations	928	928	927	928

Notes: in the top panel, the dependent variable is the log median hourly wage, by 6-digit SOC occupation in the 2014-2016 Occupational Employment Statistics. The independent variable is the log median salary, by 6-digit SOC occupation in Burning Glass, for each salary type and pay frequency, for 2010-2016. The regression is weighted least squares, weighted by 6-digit SOC occupation employment share in the OES.

In the bottom panel, the dependent variable is average weekly earnings by CBSA, from the 2010-2016 QCEW. The independent variable is the median salary by CBSA, pay frequency and salary type, from the 2010-2016 Burning Glass data. The regression is weighted least squares, weighted by CBSA employment in the QCEW.

Robust standard errors are in parentheses. One, two and three asterisks denote significance at the 5, 1 and 0.1 percent levels, respectively.

Table 3: The Length of Time Between Vacancies Posted by a Job

Length of Time Between Vacancies Posted by a Job (Quarters)	Share of Vacancies in the Sample
1	0.6
2	0.18
3	0.08
4	0.05
5	0.03
6	0.02
7	0.01
8	0.01

Notes: the first column is length of time, in quarters, between vacancies posted by a job, in Burning Glass for 2010-2016. The sample is the same as in Table 4 in the main text. The second column is the share of vacancies in the sample, for each length of time between vacancies. For example, the first row shows that for 60% of vacancies, the job posted a vacancy in the previous quarter.

Table 4: Annual Job-Level Statistics On Wage for New Hires

	Unweighted	OES Weights	QCEW Weights	High Wage Jobs
Probability of Job-Level Wage Change	0.405	0.418	0.402	0.418
Probability of Job-Level Wage Decrease	0.088	0.095	0.09	0.087
Probability of Job-Level Wage Increase	0.304	0.305	0.3	0.31
Implied Duration for which Job-Level Wages Are Unchanged (Years)	1.841	1.836	1.875	1.841

Notes: a job is an establishment by region by job title by salary type by pay frequency observation. The wage for new hires is averaged within each job-year. The sample is the 2010-2016 Burning Glass data. We estimate the probability of job-level wage change using a similar method to Klenow and Kryvtsov (2008) and 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 column (2), we reweight to target the distribution of jobs at the 6 digit SOC level from the 2014-2016 OES. In column (3) we reweight to target the distribution of employment across states from the 2010 QCEW. In column (4) we drop jobs in the bottom quartile of the wage distribution.

Table 5: Cyclicality of the Probability of Quarterly Wage Change for New Hires

Dependent Variables:	Quarterly Probability of Wage Change		Quarterly Probability of Wage Increase		Quarterly Probability of Wage Decrease	
Independent Variable:						
Change in Quarterly	-0.0255*	-0.0326*	-0.0164 +	-0.0267*	-0.00910*	-0.00596*
Unemployment	(0.00984)	(0.0142)	(0.00853)	(0.0132)	(0.00353)	(0.00257)
QCEW Weights	Y	N	Y	N	Y	N
Number of observations	1404	1404	1404	1404	1404	1404
State Clusters	52	52	52	52	52	52

Notes: the probability of a wage change for a new match is the share of vacancies for which the wage changes at the job level, in each state-quarter, from the 2010-2016 Burning Glass data. The probability of increase and decrease is defined in the same way. Wages for new hires are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. Regressions are weighted either by state employment share from the QCEW or the share of vacancies from Burning Glass.

Table 7: Job-Level Growth in Wage for New Hires and Industry Labor Demand Growth

Panel A:	Quarterly Job-Level Growth in Wage for New Hires			
	(1)	(2)	(3)	(4)
$\Delta \log\left(\text{employment}_{it}\right)$	-0.00416	-0.00188	-0.00632	0.00571**
	(0.00301)	(0.00321)	(0.00348)	(0.00206)
$\Delta \log\left(\text{employment}_{it}\right)$	0.0180***	0.0157***	0.0249***	
$\times I\left(\Delta \log\left(\text{employment}_{it}\right) > 0\right)$	(0.00425)	(0.00440)	(0.00525)	
Time Effects	Y	Y	Y	Y
Industry Trend	N	Y	N	N
Seasonally Adjusted	N	N	Y	N
Number of observations	791270	791269	791270	791270
Industry clusters	75	75	75	75

Panel B:	Annual Job-Level Growth in Wage for New Hires		
	(1)	(2)	(3)
$\Delta \log\left(\text{labor productivity}_{it}\right)$	-0.126	-0.122	-0.0125
	(0.0693)	(0.0921)	(0.0465)
$\Delta \log\left(\text{labor productivity}_{it}\right)$	0.210 +	0.244 +	
$\times I\left(\Delta \log\left(\text{labor productivity}_{it}\right) > 0\right)$	(0.108)	(0.137)	
Time Effects	Y	Y	Y
Industry Trend	N	Y	N
Number of observations	135977	135976	135977
Industry clusters	49	49	49

Notes: in Panel A, the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the growth in industry-quarter employment from the 2010-2016 Current Employment Statistics, in percentage points, at the 3 digit NAICS level. We regress quarterly job-level wage growth on quarterly industry employment growth, and interact employment growth with an indicator variable for whether employment growth is positive, in all columns but the last. Wage growth is trimmed at the 1st and 99th percentiles. In Panel B, the dependent variable is annual percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each year. The independent variable is the growth in industry-year labor productivity from the 2010-2016 BLS multifactor productivity industry accounts, in percentage points, at the 3 digit NAICS level. Labor productivity is defined as real value added per hour worked. We regress annual job-level wage growth on annual industry labor productivity growth, and interact labor productivity growth with an indicator variable for whether labor productivity growth is positive, in all columns but the last. Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by industry. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively.

Table 6: Regression of Change in Share of Jobs Posting Wage Ranges on Unemployment Changes

Dependent Variable:	Quarterly Change in Share of Jobs Posting Ranges			
	(1)	(2)	(3)	(4)
Independent Variable:				
ΔU_{st}	1.37	1.38	1.35	1.33
	(0.764)	(0.787)	(1.11)	(1.12)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-0.345	-0.319	-0.267	-0.181
	(1.23)	(1.37)	(1.77)	(1.90)
Time Effect	Y	Y	Y	Y
State Trend	N	N	Y	Y
Vacancy Weighted	Y	N	Y	N
Employment Weighted	N	Y	N	Y
N	1377	1377	1377	1377

Notes: the dependent variable is quarterly change in the share of vacancies within each state that post wage ranges, from the 2010-2016 Burning Glass data, multiplied by 100. The sample is the same as Table 5 from the main text. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. We project positive and negative unemployment changes on positive and negative employment growth changes. Column (1) and (3) weight by average state vacancies, columns (2) and (4) weight by average state employment over 2010-2016. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 8: Job-Level Growth in Wage for New Hires and State-Industry Employment Growth

Dependent Variable:	Quarterly Job-Level Growth in Wage for New Hires			
	State by 2 digit Industry		State by 3 digit Industry	
$\Delta \log(\text{employment}_{ist})$	-0.00313**	-0.00248*	-0.00276*	-0.00196
	(0.00118)	(0.00122)	(0.00131)	(0.00126)
$\Delta \log(\text{employment}_{ist})$	0.0147***	0.0125***	0.0115***	0.00958***
$\times I(\Delta \log(\text{employment}_{ist}) > 0)$	(0.00193)	(0.00199)	(0.00193)	(0.00186)
Time Effects	Y	Y	Y	Y
State-Time Effects	Y	Y	Y	Y
Industry-State Effects	N	Y	N	Y
Number of observations	1172426	1172418	1030536	1030354

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the growth in industry-state quarterly employment from the 2010-2016 Quarterly Census of Employment and Wages, in percentage points. The first two columns are at the 2 digit NAICS level, the last two columns at the 3 digit NAICS level. We regress quarterly job-level wage growth on quarterly state-industry employment growth, and interact employment growth with an indicator variable for whether employment growth is positive. Wage growth is trimmed at the 1st and 99th percentiles.

Table 9: Job-Level Growth in Wage for New Hires and City Employment Growth

	Quarterly Job-Level Growth in Wage for New Hires					
	Nominal			Real		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log\left(\text{employment}_{mt}\right)$	0.118	0.120	-0.679*	0.172**	0.172**	-0.900*
	(0.0601)	(0.0684)	(0.317)	(0.0620)	(0.0562)	(0.333)
$\Delta \log\left(\text{employment}_{mt}\right)$	0.225**	0.227*	1.098**	0.264***	0.270**	1.484***
$\times I\left(\Delta \log\left(\text{employment}_{mt}\right) > 0\right)$	(0.0708)	(0.0832)	(0.314)	(0.0690)	(0.0762)	(0.335)
Job-Level Difference	Y	Y	Y	Y	Y	Y
Time Effects	Y	Y	Y	Y	Y	Y
City Trend	N	Y	N	N	Y	N
Seasonally Adjusted	N	N	Y	N	N	Y
Number of observations	581862	581862	581862	580713	580713	580713

Notes: in Panel A, the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the growth in city-quarter employment from the 2010-2016 State and Area Employment, in percentage points, at the MSA level. We regress quarterly job-level wage growth on quarterly city employment growth, and interact employment growth with an indicator variable for whether employment growth is positive. Wage growth is trimmed at the 1st and 99th percentiles. Real wages are deflated by semiannual city prices, excluding shelter, from the Consumer Price Index.

A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by industry. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively.

Table 10: Downward Wage Rigidity by Occupation group

Dependent Variable:	Quarterly Job-Level Growth in Wage for New Hires				
Occupation Group:	Management	Services	Sales	Construction	Production
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.177**	-1.410***	-0.983*	-1.043*	-1.552***
	(0.348)	(0.310)	(0.447)	(0.433)	(0.321)
Number of Observations	568307	195274	342738	75637	329647

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. We estimate the regression separately for every broad occupation group, at the 1 digit SOC code level. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. We project positive and negative unemployment changes on positive and negative employment growth changes, and report the coefficient on the interaction term.Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 11: Downward Wage Rigidity Across the Wage and Establishment Size Distribution

Dependent variable:	Quarterly Job-Level Growth in Wage for New Hires			
Quartile of the Wage Distribution	1st	2nd	3rd	4th
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.078***	-0.889**	-1.318***	-1.460**
	(0.243)	(0.267)	(0.325)	(0.434)
N	483806	327623	367982	386771
Quartile of Establishment Size	1st	2nd	3rd	4th
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.573***	-1.278***	-0.884**	-1.069**
	(0.280)	(0.328)	(0.328)	(0.382)
N	365399	461153	265379	358835

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. We project positive and negative unemployment changes on positive and negative employment growth changes.

We run the regression separately for each quartile of the wage distribution, in the top panel. The wage distribution is the occupation's rank in the distribution of median hourly wages, from the 2014-2016 Occupational Employment Statistics. We run the regression for each quartile of the establishment size distribution, in the bottom panel. The size distribution is the number of vacancies posted by the establishment between 2010 and 2016.

Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state; except for the last column, which clusters by state and quarter. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. In some specifications, we reweight to target the occupation employment distribution at the 6 digit SOC level from the 2014-2016 OES.

Table 12: Downward Wage Rigidity By Source of Vacancy

Dependent variable:	Quarterly Job-Level Growth in Wage for New Hires			
Source	Company Website	Job Board	Government	Education
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.427**	-1.409***	-0.969	-1.432*
	(0.440)	(0.375)	(1.423)	(0.668)
N	149261	513955	95650	53506
Time Effect	Y	Y	Y	Y

Notes: the dependent variable is quarterly percentage wage growth for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. We project positive and negative unemployment changes on positive and negative employment growth changes.

We run the regression separately for the largest four sources of data on which Burning Glass draws. The wage distribution is the occupation's rank in the distribution of median hourly wages, from the 2014-2016 Occupational Employment Statistics. The first column studies exclusively vacancies that post on company websites, the second column on online job boards, the third column on government websites and the last column on the websites of educational institutions.

Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. Standard errors are in parentheses, clustered by state; except for the last column, which clusters by state and quarter. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. In some specifications, we reweight to target the occupation employment distribution at the 6 digit SOC level from the 2014-2016 OES.

Table 13: Decomposition of Job-Level Wage Growth

Dependent Variable:	$I\left(\Delta w_{ijt} \neq 0\right)$		$\Delta \log w_{ijt} \vert w_{ijt} \neq 0$	
	(1)	(2)	(3)	(4)
ΔU_{st}	-0.0991	-0.0790	-0.193	-0.210
	(0.0752)	(0.0541)	(0.623)	(0.812)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-0.0486	-0.0336	-0.494*	-0.409
	(0.0259)	(0.0204)	(0.214)	(0.303)
Time Effect	Y	Y	Y	Y
State Trend	N	Y	N	Y
N	1598139	1598139	373375	373375

Notes: in the first two columns, the dependent variable is an indicator for whether wages change at the job level. In the last two columns, the outcome variable is wage growth conditional on wages changing. Wages are averaged within each job-quarter. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW, in all but the final column. We project positive and negative unemployment changes on positive and negative employment growth changes. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 14: First Stage of Quarterly State Unemployment Change on Employment Growth

Dependent Variable:	Quarterly Unemployment Change			
	(1)	(2)	(3)	(4)
Independent Variable:				
Quarterly Employment Growth	-0.215***	-0.216***	-0.262***	-0.263***
	(0.0265)	(0.0262)	(0.0157)	(0.0157)
State Difference	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y
State Trend	N	Y	N	Y
QCEW Weight	N	N	Y	Y
Number of Observations	1404	1404	1404	1404
R^2	0.599	0.631	0.637	0.663
F Statistic	66.14	67.78	277.8	282.1
State Clusters	52	52	52	52

Notes: the dependent variable is the quarterly change in state level unemployment, from the 2010-2016 LAUS, in percentage points. The independent variable is the quarterly growth in state level employment from the 2010-2016 QCEW, in percentage points. In columns (3) and (4), the regression is weighted least squares, reweighted to target average state level employment in the QCEW. Standard errors are in parentheses, clustered by state. One, two and three asterisks denote significance at the 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 15: Regression of Establishment Share in High Wage Occupations on Unemployment

	Quarterly Change in Share of Establishment Vacancies					
	in High Wage Occupations		with High Wages		in High Wage Occupations, by Broad Occupation Group	
ΔU_{st}	0.158	0.922	-0.0296	-0.0031	0.00881	-0.0728
	(0.370)	(0.642)	(0.0358)	(0.0577)	(0.338)	(0.333)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	0.167	-0.158	0.0537	0.0173	0.509	0.772
	(0.441)	(0.735)	(0.0369)	(.0648)	(0.351)	(0.341)
Establishment Difference	Y	Y	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y	Y	Y
Size Weighted	N	Y	N	Y	N	Y
N	1770257	1770257	1883361	1883361	2388716	2388716
State Clusters	52	52	52	52	52	52

Notes: In the first two columns, the dependent variable is the change in the quarterly share of establishment vacancies in high wage occupations. High wage occupations are occupations with wages above the weighted median wage, by occupation, in 2010-2016 Burning Glass. The occupations are defined at the 6 digit SOC code level, occupation wages are the median hourly wage according to the 2014-2016 Occupational Employment Statistics.

In the middle two columns, the dependent variable is the change in the quarterly share of high wage establishment vacancies. High wage vacancies are vacancies with wages above the weighted median wage within each pay frequency (e.g. hourly or annual) and salary type (e.g. total or base pay). The occupations are again at the 6 digit SOC level.

In the final two columns, the dependent variable is the change in the quarterly share of establishment vacancies in high wage occupations, within broad occupation groups. A high wage occupation within a broad occupation group, is a 6 digit SOC occupation, that is above the vacancy-weighted median hourly wage within the set of 6 digit SOC occupations belonging to the same broad occupation group. For example, CEOs (6 digit SOC code 11-1011) have above the median wage of the occupations belonging to the 1 digit SOC occupation group of Management, Business, Science, and Arts Occupations. The broad occupation groups are the set of 6 occupation groupings defined by the BLS in 2018. Size weighted denotes weighted by establishment-quarter size.

In all columns, the independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, in percentage points. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW. In the asymmetric specifications, we project positive and negative unemployment changes on positive and negative employment growth changes. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

Table 16: Regression of Log Labour Productivity on Unemployment

Dependent Variable:	Log Labour Productivity Change			
	Quarterly		Annual	
Independent Variable:				
Unemployment Change	-0.537	-0.858	-2.242	-6.450
	(1.012)	(1.790)	(1.141)	(3.286)
Time Effect	Y	Y	Y	Y
State Effect	N	Y	N	Y
Number of Observations	1377	1377	306	306
State Clusters	52	52	52	52

Notes: the dependent variable is the change log labour productivity for 2010-2016. Labour productivity is defined as gross state product from the BEA's regional economic accounts, divided by the number of employed from the QCEW. The independent variable is the change in state unemployment from the 2010-2016. We project changes in state unemployment on growth in state employment from the 2010-2016 QCEW. Standard errors are in parentheses, clustered by state. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico. The regressions are weighted by 2010 state level employment from the QCEW.

C Additional Empirics

C.1 Testing the Constant Hazard Assumption

In this subsection we test the constant hazard assumption of our model. In particular, we estimate the hazard model in footnote 18 separately for each value of γ_{it}, the gap in quarters between the wage at t and wage in the previous vacancy that was posted, for job i. If the estimated hazard is similar for all values of γ_{it}, then the constant hazard assumption is satisfied. The following table estimates the hazard for values of γ_{it} between 1 and 8 quarters. The estimated value of the hazard is stable for values of γ_{it} in this range.

Time Since Last Post	Estimated Hazard
1	0.16
2	0.16
3	0.17
4	0.17
5	0.17
6	0.16
7	0.16
8	0.14

C.2 The Cyclicality of Education Requirements

This section studies the cyclicality of education requirements.

Burning Glass reports whether a vacancy requires education in one of five categories, namely: a high school degree, associate's degree, bachelor's degree, masters degree or PhD. We assign to each education requirement the average number of years of education required, which is, respectively, 12, 14, 16, 18 and 21 years.

We construct measures of whether a job changes its education requirement. For each job, with non-missing education requirements, we study the change in the number of years of education required, between consecutive vacancies posted by the job. In our sample of jobs which post multiple vacancies, roughly half of the vacancies post education

We estimate the regression

$$\Delta\text{years of education}_{jst} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{jst}.$$

This regression is identical to our benchmark regression equation (1), but replaces as the outcome variable the change in the number of years of education required, instead of the growth in the job level wage, as in the benchmark regression. When a job posts multiple vacancies in the quarter with non missing education requirements, we take the mean years of education required.

The following table reports the results. Across all specifications, education requirements respond to neither increases nor decreases in unemployment.

Dependent Variable:	Quarterly Change in Years of Education			
	(1)	(2)	(3)	(4)
Independent Variable:				
ΔU_{st}	-0.00829	-0.00673	0.00505	0.000352
	(0.0136)	(0.0151)	(0.0161)	(0.00276)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	0.00883	0.00737	-0.00486	
	(0.0130)	(0.0146)	(0.0158)	
Job-Level Difference	Y	Y	Y	Y
Time Effect	Y	Y	Y	Y
State Trend	N	Y	N	N
OES Weight	N	N	Y	N
N				

C.3 Wage Growth After Being Hired

This subsection studies wage growth after workers are hired. First, we develop a measure of wage growth after workers are hired. Second, we develop an empirical proxy for this measure using worker level survey data. Third, we estimate the cyclicality of wage growth after workers are hired. We find that during contractions, wages grow faster for workers after they are hired.

First, we construct a measure of wage growth after workers are hired. Kudlyak (2014) shows that the present value of wages matters for job creation in canonical labor search models. The present value of wages PDV_t is

$$PDV_t = E_t \sum_{j=0}^{\infty} \left[\beta(1-s)\right]^j w_{t,t+j}$$

where $w_{t,t+j}$ is the wage in period $t+j$ for a worker hired in period t, s is the job separation rate and β is the discount factor. Rearranging the logarithm of this expression yields

$$\log PDV_t = \underbrace{\log\left(E_t \sum_{j=0}^{\infty} \left[\beta(1-s)\right]^j \frac{w_{t,t+j}}{w_{t,t}}\right)}_{\text{wage growth after being hired}} + \overbrace{\log w_{t,t}}^{\text{wages for new hires}}. \tag{18}$$

The first term is wage growth for workers, after they are hired, relative to the initial wage when they are hired. The second term is the wage paid to workers when they are hired. The main body is devoted to estimating the cyclicality of the second term, wages for new hires. In this subsection, we estimate the cyclicality of the second term, wage growth after being hired. In effect we are estimating the component of the "user cost of labor" that is not explained by the wage for new hires.

Both components of the right hand side equation (18) matter for the present value of wages. Suppose that, as we find in the main text, wages for new hires do not fall during contractions. Subsequent wage growth for new hires might still be lower. Then, the present value of wages would fall during contractions. So, firms' cost of labor could still fall.

We emphasized in the main text that estimates of the cyclicality of $\log w_{t,t}$ may be confounded by job composition. But estimates of $E_t \sum_{j=0}^{\infty} \left[\beta(1-s)\right]^j \frac{w_{t,t+j}}{w_{t,t}}$ will not be confounded by job composition. The reason is, we are measuring wage growth *within the same job* after a worker is hired into that job. So, the composition of jobs is held fixed.

Now, we explain how we construct an empirical proxy of $E_t \sum_{j=0}^{\infty} \left[\beta(1-s)\right]^j \frac{w_{t,t+j}}{w_{t,t}}$, wage growth after being hired. Our methods follow almost exactly the state of the art in Kudlyak (2014) and Basu and House (2016). We use annual data from the National Longitudinal Survey of Youth between 1979 and 2013. We focus only on men. We reweight to make the data representative of the US population. Our measure of the wage is the "hourly rate of pay" variable. We construct real wages using the deflator for the nonfarm business sector. We then construct empirical measures of $w_{t,t+j}$ in exactly the same way as Kudlyak (2014) and Basu and House (2016).

We then study the cyclicality of wage growth after being hired. We run the regression

$$\Delta \log\left(\text{wage growth after being hired}_t\right) = \alpha + \beta \Delta U_t + \gamma \Delta U_t \times I\left(\Delta U_t < 0\right) + \delta I\left(\Delta U_t < 0\right) + \varepsilon_t$$

where U_t is annual unemployment from the Bureau of Labor Statistics.

The table at the end of the subsection reports the results. The results suggest that during contractions, wage growth after being hired increases. Still, the results are imprecise. In the first row, the coefficient on ΔU_t is positive. So, when unemployment rises by a percentage point, then wage growth after being hired increases by 0.7%. In the second row, the coefficient on $\Delta U_t \times I\left(\Delta U_t < 0\right)$ is negative and of a similar magnitude to the coefficient on ΔU_t. When unemployment falls by a percentage point, wage growth for new hires increases by the sum of the coefficients in the two rows of the table. So, wage growth for new hires increases by $0.76\% - 0.53\% = 0.23\%$ after a percentage point fall in unemployment.

	$\Delta \log\left(\text{wage growth after being hired}_t\right)$
$\Delta\text{Unemployment}_t$	0.763
	(1.19)
$\Delta\text{Unemployment}_t \times$	-0.529
$I\left(\Delta\text{Unemployment}_t < 0\right)$	(-0.73)
N	28

C.4 The Cyclicality of the Difference between Measures of the Wage for New Hires

This section asks whether the difference between measures of the wage for new hires in Burning Glass, and alternative measures of the wage for new hires from survey data, correlate with unemployment fluctuations. This step is important because Burning Glass measures the wage posted on vacancies, not the realized wage paid to new hires. So, if the gap between these wages are cyclical, then our estimates of downward wage rigidity in the main text are hard to interpret.

As an outcome variable, we study the gap between the average wage for new hires in Burning Glass, and the average wage for new hires from the Current Population Survey, at the state by quarter level. We explain how we construct average wages for new hires in the Current Population Survey in section 2.1. We difference this variable and then regress it on changes in quarterly unemployment. Our regression equation is

$$\Delta \text{wage gap}_{st} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{st}$$

where γ_t is a time fixed effect, ΔU_{st} is the change in quarterly state unemployment, and $\Delta \text{wage gap}_{st}$ is the difference between wages for new hires in Burning Glass and another data source. As in the main text, we instrument for changes in unemployment with employment growth.

The following table reports the results. Though noisy, the data do not suggest a significant comovement between the wage gap and unemployment changes.

Dependent Variable:	$\Delta \text{wage gap}_{st}$			
	(1)	(2)	(3)	(4)
ΔU_{st}	-8.514	-9.916	-6.189	-8.295
	(5.915)	(6.450)	(6.070)	(7.681)
$\Delta U_{st} \times I\left(\Delta U_{st} < 0\right)$	9.939	11.66	8.965	11.72
	(6.669)	(7.737)	(7.780)	(9.863)
N	1377	1377	1377	1377

C.5 Can Establishment Level Hiring Offset Downward Wage Rigidity?

One potential concern is that establishments alter the jobs into which they hire workers, in a way that offsets downward rigidity at the job level. Granted, since establishment- and job-level wages display a similar degree of downward rigidity, this concern does not seem to matter in practice. Nevertheless, we explain the concern and how we deal with it.

Consider a simple example. Suppose that in the Starbucks establishment, wages are downwardly rigid for "senior baristas" and "junior baristas". During expansions, Starbucks hires higher wage senior baristas. During contractions, Starbucks hires lower wage junior baristas. Either way, newly hired workers brew coffee. The wage for new hires falls despite downward rigidity at job level, without any effect on the output of the Starbucks. More generally, establishments could avoid wage rigidity at the job title level. During booms, establishments could hire in high wage jobs; and during busts, hire in low wage jobs. Then the wage faced by the establishment might fall during contractions, and partially offset the effect of downward wage rigidity.

This concern supposes that establishments can easily substitute between high and low wage workers. In practice, low and high wage jobs might be very different, preventing such substitution.

We test whether establishments circumvent job level wage rigidity in this manner, by asking whether establishments increase their hiring in low wage jobs during contractions. For each establishment, and in each quarter, we calculate the share of high wage vacancies, with three methods. First, we calculate the share of establishment-quarter vacancies that are above the weighted median wage in Burning Glass. Second, we calculate the share of vacancies that are above the median 6 digit SOC occupation wage, that is the share of vacancies in high wage occupations. Third, we calculate the share of vacancies in high wage occupations, within each establishment and broad occupation group. A broad occupation group is at the 1 digit SOC level. This third method contemplates that establishments might substitute jobs differently, depending on the broad occupation group to which the job belongs.

We regress the quarterly change in the high wage establishment share, from these three measures, on the change in quarter-by-state unemployment. The regression is identical to regression equation (2)—but for the outcome variable, which is the quarterly change in the high wage establishment share.

Appendix Table (15) presents the results. Row (1) of column (1) shows that when unemployment rises by one percentage point, the share of high wage occupations in the establishment rises by 0.1 percentage points. This coefficient is small in magnitude, not statistically significant. The sign of the coefficient suggests that, if anything, establishments raise the high wage share of jobs during recessions. Row (2) of column (1) shows that the share of establishment high wage vacancies does not respond significantly differently to contractions versus expansions in unemployment. Thus the establishment share of vacancies cannot be moving in a way that offsets the asymmetric response of wages to contractions versus expansions. The results are similar with the other two methods for calculating establishments' high wage shares. So, the mix of jobs into which establishments are hiring cannot be moving to offset the downward constraint on wage setting at the job level.

C.6 Wage Ranges

Roughly half of the wage data posts a range of wages, instead of a point wage. In most specifications in the main text, we take the mean wage for jobs that post a range of wages.

Here, we show that workers in occupations with a high share of jobs that post ranges, instead of point wages, do not have more cyclical wages. Instead, dynamics in the wage for new hires are similar for jobs that post either point wages or ranges. Wage ranges do not create an additional source of wage flexibility.

To do this, we study the wage for newly hired workers in the CPS. For each worker, we classify their 3 digit SOC occupations in the CPS, as either likely to post a range, or likely to post a point wage. We classify an occupation as likely to post a wage, if the occupation has an above median share of vacancies posting a point wage in Burning Glass data.

We regress log wages for newly hired workers on quarterly state unemployment. We also interact state unemployment with an indicator for whether the worker's occupation is likely to post a wage range. If this indicator is significant, then occupations that tend to post wage ranges have different wage dynamics from occupations that tend to post point wages.

Table 17 reports the results. Occupations that are likely to post a range instead of a point have wages that are *less* responsive to regional unemployment fluctuations. The coefficient is not significant. Therefore the distinction between posting a range or posting a point wage is unlikely to matter for understanding wage cyclicality.

Table 17: Wage Cyclicality in Occupations with High vs. Low Share Posting Wage Ranges

Dependent Variable:	Log Wage, CPS, Newly Hired Workers
Independent Variables:	
Quarterly Unemployment	-1.019
	(1.11)
Quarterly Unemployment ×	1.120
High Share Posting Wage Ranges	(0.82)
Annual Unemployment	-1.131
	(1.19)
Annual Unemployment ×	1.174
High Share Posting Wage Ranges	(0.83)
Time Effect	Y
State Effect	Y
Number of Observations	67327

Notes: In Burning Glass, we classify three digit SOC occupations with an above-median and below-median share posting ranges instead of point wages. We link these occupations to the same three digit SOC occupations in the CPS. In the CPS, we denote three digit SOC occupations with above-median shares, as measured in the Burning Glass data, as having a high share posting wage ranges, and otherwise a low share. The dependent variable is usual hourly earnings, including overtime, for hourly and non-hourly workers, for new hires, which we construct following the "wage 4" series from CEPR. The wage is from the 2012-2017 CPS Merged Outgoing Rotation Group. We identify new hires by longitudinally linking workers to the previous three monthly survey waves, and isolating workers transitioning into new jobs. The independent variable is unemployment from the 2010-2016 LAUS. We project unemployment onto log employment from the QCEW. One, two and three asterisks denote significance at the 5, 1 and 0.1 percent levels, respectively. Standard errors are clustered by state.

C.7 Wage Flexibility Upward Over Time

We estimate the regression

$$\Delta \log w_{jst} = \alpha_y + \gamma_t + \beta_y \Delta U_{st} + \varepsilon_{jst}, \tag{19}$$

for $y \in \{2010, \ldots, 2016\}$, and again restrict the sample to observations where $\Delta U_{st} < 0$. That is, we estimate the regression for every year y. β_y measures the sensitivity of wage growth to falls in unemployment, estimated separately for every year y. A more negative number indicates that wage growth is more sensitive to falls in unemployment. Hence β_y is estimated using state-by-quarter panel variation, within each year y. As before, we project unemployment changes on employment growth from the QCEW to deal with measurement error.

Table 18 reports the results. During the early part of the sample period, the wage for new hires does not rise after falls in unemployment. At the end of the period, when labor markets are tighter, the wage for new hires rises strongly as unemployment falls. The rich variation also underscores the benefit of our dataset. We can precisely estimate wage cyclicality regressions on a state-quarter panel, separately for every year in our panel.

Table 18: Regression of Wage Growth on State Unemployment Declines

Dependent Variable:	Job-Level Growth in Wage for New Hires			Establishment-Level Growth in Wage for New Hires
	(1)	(2)	(3)	(4)
Independent Variables:				
ΔU_{st}	-0.208	-0.219	-0.362*	0.0261
	(0.167)	(0.180)	(0.156)	(0.493)
$\Delta U_{st} \times I(\text{Year} = 2011)$	-0.0330	0.0305	0.0416	-0.691
	(0.240)	(0.218)	(0.209)	(0.786)
$\Delta U_{st} \times I(\text{Year} = 2012)$	-0.462*	-0.379	-0.312	-0.850
	(0.221)	(0.238)	(0.230)	(0.720)
$\Delta U_{st} \times I(\text{Year} = 2013)$	-0.415	-0.360	-0.278	-0.964
	(0.244)	(0.232)	(0.216)	(0.506)
$\Delta U_{st} \times I(\text{Year} = 2014)$	-0.451*	-0.406*	-0.314	-1.045
	(0.193)	(0.185)	(0.174)	(0.540)
$\Delta U_{st} \times I(\text{Year} = 2015)$	-1.014***	-0.935***	-0.812***	-1.216*
	(0.184)	(0.190)	(0.173)	(0.516)
$\Delta U_{st} \times I(\text{Year} = 2016)$	-1.746***	-1.664***	-1.524***	-1.253*
	(0.176)	(0.184)	(0.174)	(0.474)
Time Effects	Y	Y	Y	Y
State Effects	N	Y	N	N
State × Quarter-of-Year Effects	N	N	Y	N
Number of Observations	1090035	1089914	1090035	1279369
State Clusters	52	52	52	52

Notes: we estimate the regression

$$\Delta \log w_{jst} = \sum_{y \in \{2010,\ldots,2016\}} \alpha_y + \gamma_t + \sum_{y \in \{2010,\ldots,2016\}} \beta_y I\left(\text{year} = y\right) \Delta U_{st} + \varepsilon_{jst}.$$

The dependent variable in the first three columns is quarterly job-level wage growth, in percentage points, for new hires, from the 2010-2016 Burning Glass data. Wages are averaged within each job-quarter. The dependent variable in the last column is quarterly establishment-level wage growth, in percentage points, for new hires. The independent variable is the change in state-quarter unemployment from the 2010-2016 LAUS, also in percentage points. We restrict the sample only to observations for which $\Delta U_{st} < 0$. We project unemployment changes onto state-quarter employment growth from the 2010-2016 QCEW, and both unemployment changes and employment are interacted with dummy variables for each year. Wage growth is trimmed at the 1st and 99th percentiles. A job is an employer by location by pay frequency by salary type by job title unit. An establishment is an employer by location by pay frequency by salary type unit. Standard errors are in parentheses, clustered by state. A plus sign, one, two and three asterisks denote significance at the 10, 5, 1 and 0.1 percent levels, respectively. The sample is vacancies in the 50 states, plus the District of Columbia and Puerto Rico.

D Proofs

D.1 Proof of Proposition 1

Summing regression equation (5) over i yields

$$\sum_i \nu_{ist}\Delta \log w_{ist} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta_{\text{Benchmark}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{st} \tag{20}$$

where $\varepsilon_{st} = \sum_i \nu_{ist}\varepsilon_{ist}$. We can substitute equation (4) into equation (6) to rewrite the regression that uses average wages as

$$\sum_i \nu_{ist}\Delta \log w_{ist} + \sum_i \log w_{ist}\Delta \nu_{ist} = \bar{\alpha} + \bar{\gamma}_t + \bar{\beta}\Delta U_{st} + \delta_{\text{Average}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \bar{\varepsilon}_{st}. \tag{21}$$

For notational simplicity, we can rewrite equation (20) as

$$y_{st} = \mathbf{x}_{st}^{'}\mathbf{b} + \varepsilon_{st}$$

and equation (21) as

$$y_{st} + u_{st} = \mathbf{x}_{st}^{'}\bar{\mathbf{b}} + \bar{\varepsilon}_{st}$$

where

$$y_{st} \equiv \sum_i \nu_{ist}\Delta \log w_{ist}$$

$$u_{st} \equiv \sum_i \log w_{ist}\Delta \nu_{ist}.$$

$\mathbf{x}_{st}^{'}\mathbf{b}$ and $\mathbf{x}_{st}^{'}\bar{\mathbf{b}}$ collect the covariates and coefficients in regressions 20) and (21) respectively. The OLS estimator of $\mathbf{b}$, which we term $\hat{\mathbf{b}}$, is

$$\hat{\mathbf{b}} = \left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1}\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}y_{st}\right).$$

The variance of $\hat{\mathbf{b}}$ conditional on $\mathbf{x}_{st}$ is

$$\begin{aligned} V\left[\hat{\mathbf{b}}|\mathbf{x}_{st}\right] &= V\left[\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1}\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}y_{st}\right)|\mathbf{x}_{st}\right] \\ &= \left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1}\frac{1}{(ST)^2}V\left[\left(\sum_{s,t}\mathbf{x}_{st}y_{st}\right)|\mathbf{x}_{st}\right]\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1} \end{aligned}$$

The OLS estimator of $\bar{\mathbf{b}}$, which we term $\hat{\bar{\mathbf{b}}}$, is

$$\hat{\bar{\mathbf{b}}} = \left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1}\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\left(y_{st}+u_{st}\right)\right).$$

Then the variance of $\hat{\bar{\mathbf{b}}}$ conditional on $\mathbf{x}_{st}$ is

$$V\left[\hat{\bar{\mathbf{b}}}|\mathbf{x}_{st}\right] = V\left[\hat{\mathbf{b}}|\mathbf{x}_{st}\right] + \left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1}\frac{1}{(ST)^2}V\left[\sum_{s,t}\mathbf{x}_{st}u_{st}|\mathbf{x}_{st}\right]\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}_{st}'\right)^{-1} \tag{22}$$

The second term in equation (22) is a matrix with strictly positive entries on its leading diagonal for $S, T < \infty$. Hence every entry on the leading diagonal of $V\left[\hat{\bar{\mathbf{b}}}|\mathbf{x}_{st}\right]$ is greater than the corresponding entry on the leading diagonal of $V\left[\hat{\mathbf{b}}|\mathbf{x}_{st}\right]$.

D.2 Derivation of Equation (16)

Throughout, $\bar{x}$ is the steady state value of x_t.

Let y_t follow an AR(1) process with first order autocorrelation ρ and mean value 1. We derive equation (16) assuming that entry is positive, so $v_{it} > 0$ for all t. Given our assumptions on wage setting, entry will indeed be positive for both job types in a neighborhood of the steady state. Substituting the free entry condition (15) into the Bellman equation for jobs (14) yields

$$\begin{aligned} J_{it,t} &= \mathbb{E}_t\sum_{j=0}^{\infty}\left[\beta(1-s)(1-\omega)\right]^j\left(y_{t+j}-w_{it}\right) \\ &= \frac{y_t-1}{1-\rho\beta(1-\omega)(1-s)} + \frac{1}{1-\beta(1-\omega)(1-s)} - \frac{w_{it}}{1-\beta(1-\omega)(1-s)} \end{aligned}$$

Substituting the free entry condition (15) into the Bellman equation for vacancies (13) yields

$$J_{it,t} = \frac{c}{q(\theta_{it})}.$$

Equating these last two equations yields

$$\frac{y_t-1}{1-\rho\beta(1-\omega)(1-s)} + \frac{1}{1-\beta(1-\omega)(1-s)} - \frac{w_{it}}{1-\beta(1-\omega)(1-s)} = \frac{c}{q(\theta_{it})} \tag{23}$$

$$\implies \frac{\Delta\log\theta_{it}}{\Delta\log y_t} = \frac{1}{\alpha}\frac{1}{1-\bar{w}_i}\left(\frac{1-\beta(1-\tau)}{1-\rho\beta(1-\tau)} - \bar{w}_i\frac{\Delta\log w_{it}}{\Delta\log y_t}\right) \tag{24}$$

where

$$\tau \equiv 1-(1-\omega)(1-s).$$

and I use a first order approximation around the steady state in the third line. We can also define the steady state tightness in each job, using equations (8), (9) and (23), as

$$\frac{1}{1-\beta(1-s)(1-\omega)} - \frac{\bar{w}_i}{1-\beta(1-s)(1-\omega)} = \frac{c}{q(\bar{\theta}_i)}$$

$$\implies \frac{1}{1-\beta(1-s)(1-\omega)} - \frac{\bar{w}_i}{1-\beta(1-s)(1-\omega)} = c\Psi\bar{\theta}_i^{\alpha}. \tag{25}$$

We have from equation (12)

$$l_{it} = \frac{1}{2} - \left(1-f(\theta_{it})\right)\left(\frac{1}{2} - (1-\omega)(1-s)\,l_{i,t-1}\right).$$

We make the approximation that given θ_{it}, $l_{it} = l_{i,t-1}$, i.e. flows into and out of employment are approximately equal conditional on current market tightness. Shimer (2005) verifies that this approximation is accurate at quarterly frequency. Simplifying equation (12) with this approximation yields

$$l_{i,t-1} = \frac{1}{2} - \left(1-f(\theta_{it})\right)\left(\frac{1}{2} - (1-\omega)(1-s)\,l_{i,t-1}\right)$$
$$\implies l_{i,t-1} = \frac{1}{2}\frac{1}{\frac{\tau}{f(\theta_{it})} + (1-\tau)}.$$

We also have from equation (10)

$$\begin{aligned} u_{it} &= \frac{1}{2} - (1-\omega)(1-s)\,l_{i,t-1} \\ &= \frac{1}{2}\left(\frac{\tau}{\tau+\theta_{it}q(\theta_{it})(1-\tau)}\right) \end{aligned}$$

$$\begin{aligned} \implies \frac{\Delta u_{it}}{\Delta\theta_{it}} &= -\frac{1}{2}\frac{\tau}{\left[\tau+\theta_{it}q(\theta_{it})(1-\tau)\right]^2}(1-\tau)\left[q(\theta_{it})+\theta_{it}q'(\theta_{it})\right] \\ &= -u_{it}q(\theta_{it})(1-\alpha)\frac{(1-\tau)}{\left[\tau+f(\theta_{it})(1-\tau)\right]} \end{aligned}$$

$$\begin{aligned} \implies \frac{\Delta u_{it}}{\Delta\theta_{it}}\frac{\theta_{it}}{u_{it}} &= -(1-\alpha)\frac{(1-\tau)\theta_{it}q(\theta_{it})}{\left[\tau+f(\theta_{it})(1-\tau)\right]} \\ &= -(1-\alpha)(1-2u_{it}) \end{aligned}$$

$$\implies \frac{\Delta\log u_{it}}{\Delta\log\theta_{it}} = -(1-\alpha)(1-2u_{it}). \tag{26}$$

Therefore in a neighborhood of the steady state

$$\begin{aligned}\frac{\Delta\log u_{it}}{\Delta\log y_t} &= \frac{\Delta\log u_{it}}{\Delta\log\theta_{it}}\frac{\Delta\log\theta_{it}}{\Delta\log y_t} \\ &= -\frac{(1-\alpha)(1-2\bar{u}_i)}{\alpha}\frac{1}{1-\bar{w}_i}\left(\frac{1-\beta(1-\tau)}{1-\rho\beta(1-\tau)} - \bar{w}_i\frac{\Delta\log w_{it}}{\Delta\log y_t}\right),\end{aligned} \tag{27}$$

where the second line uses equations (26) and (24). Since $u_t = u_{Ht} + u_{Lt}$, we have

$$\begin{aligned}\frac{\Delta\log u_t}{\Delta\log y_t} &= \frac{\Delta\log(u_{Ht}+u_{Lt})}{\Delta\log y_t} \\ &= -A + B\frac{\mu\Delta\log w_{Ht} + (1-\mu)\Delta\log w_{Lt}}{\Delta\log y_t}\end{aligned}$$

where we use equation (27) and

$$A \equiv \frac{1}{\bar{u}_H+\bar{u}_L}\sum_{i=H,L}\frac{\bar{u}_i(1-\alpha)(1-2\bar{u}_i)}{\alpha}\frac{1}{1-\bar{w}_i}\frac{1-\beta(1-\tau)}{1-\rho\beta(1-\tau)}$$

$$B \equiv \frac{\sum_{i=H,L}\frac{\bar{u}_i(1-\alpha)(1-2\bar{u}_i)}{\alpha}\frac{1}{1-\bar{w}_i}\bar{w}_i}{\bar{u}_H+\bar{u}_L}$$

$$\mu \equiv \frac{\frac{\bar{u}_H(1-\alpha)(1-2\bar{u}_H)}{\alpha}\frac{1}{1-\bar{w}_H}\bar{w}_H}{\sum_{i=H,L}\frac{\bar{u}_i(1-\alpha)(1-2\bar{u}_i)}{\alpha}\frac{1}{1-\bar{w}_i}\bar{w}_i}$$

where $\bar{w}_i = \phi_i$.

D.3 Derivation of Equation (17)

For $\Delta\log y_t > 0$ we have from equations (8) and (9)

$$w_{it} = \phi_i y_t^{\gamma}$$

$$\implies \Delta\log w_{it} = \gamma\Delta\log y_t$$

and so from equation (**??**)

$$\begin{aligned}\frac{\Delta\log u_t}{\Delta\log y_t} &= -A + B\frac{\mu\gamma\Delta\log y_t + (1-\mu)\gamma\Delta\log y_t}{\Delta\log y_t} \\ &= -A + B\gamma.\end{aligned}$$

Analogously, for $\Delta \log y_t < 0$ we have

$$\frac{\Delta \log u_t}{\Delta \log y_t} = -A.$$

E Model with Heterogeneous Jobs

The goal of this model is to show that wage rigidity at the job level is particularly important for unemployment fluctuations.

We study a standard Diamond-Mortensen-Pissarides model, augmented with heterogeneous jobs. So, there is a concrete definition of the job level wage.

We model jobs differently from the main text. In this model, we study a continuum of firms, with heterogeneous productivity and decreasing returns to scale. Each firm corresponds to a job. This model is a much simplified version of Elsby and Michaels (2013) and also closely relates to Michaillat (2012). Amongst the many simplifications relative to Elsby and Michaels (2013), we assume an exogenous separation rate.

E.1 Model Setup

The model is in discrete time. Fluctuations in unemployment are determined by a random process for labor productivity z_t.

There is a unit mass of workers in the labor market. Workers are either employed, or they are unemployed and search for jobs at firms. Workers have risk neutral preferences, discount future payoffs with discount factor $\beta \in (0,1)$, and consume all their income in each period. If workers are unemployed, then they have zero consumption.

There is a unit measure of firms $i \in [0,1]$ that hire workers. Each firm has random and idiosyncratic labor productivity x_{it}. Each firm operates a decreasing returns production technology.

For notation, unless otherwise mentioned, variables indexed by i and t refer to the variable of firm i at time t. The same variable without an i subscript denotes the corresponding average of the variable across firms. The same variable without a t subscript denotes steady state values.

E.1.1 Labor Market

Our model of the frictional labor market follows exactly the standard Diamond-Mortensen-Pissarides framework. At the end of $t-1$, an exogenous share s of the $n_{i,t-1}$ workers employed at firm i separate from the firm. These workers search for new jobs at once. So, at the start of period t, u_t unemployed workers search for jobs whereby

$$u_t = 1 - (1-s)\, n_{t-1} \tag{28}$$

where $n_{t-1} = \int_0^1 n_{i,t-1} di$. Firm i posts v_{it} vacancies to hire unemployed workers.

Labor market conditions are captured by labor market tightness $\theta_t = u_t / v_t$ where $v_t = \int_0^1 v_{it} di$. The number of matches made in a period is given by $m_t = \Psi u_t^{\iota} v_t^{1-\iota}$. So, the probability that a firm fills each vacancy is given by

$$q(\theta_t) = \frac{\Psi u_t^{\iota} v_t^{1-\iota}}{v_t} = \Psi \theta_t^{-\iota}$$

where $\iota \in (0,1)$ so that q is decreasing in θ_t. When the labor market is tight, then many firms post vacancies to match with few workers. So, the probability that an individual vacancy is filled is low. The probability that an unemployed worker finds a job is

$$f(\theta_t) = \frac{\Psi u_t^{\iota} v_t^{1-\iota}}{u_t} = \Psi \theta_t^{1-\iota},$$

which is decreasing in θ_t. When the labor market is tight, the probability that an individual unemployed worker finds a job is relatively low. Vacancy posting incurs a cost c per vacancy, for each period in which the vacancy is open. There is no randomness at the firm level. Firm i hires $h_{it} \geq 0$ workers after posting $h_{it}/q(\theta_t)$ vacancies. Workers start producing output in the period that they are hired.

At the steady state of this economy, there is a so-called Beveridge curve—a positive relationship between tightness and employment. At the steady state, flows into unemployment sn and flows out of unemployment $[1-(1-s)\, n]\, f(\theta)$ are equal, so

$$sn = [1-(1-s)\, n]\, f(\theta)$$

$$\implies n = \frac{1}{(1-s) + \frac{s}{f(\theta)}}.$$

As firms post more vacancies v, tightness θ rises. Then $f(\theta)$ rises, as workers are more likely to find jobs. So, n rises.

E.1.2 Wages

Our first departure from the standard DMP model is to allow downward wage rigidity.

If a firm i successfully matches with an unemployed worker at time t, then they set a wage

w_{it}. We specify a linear wage rule

$$w_{it} = \max\left[\varphi_w + \varphi_z z_t + \varphi_x x_{it}, \frac{w_{i,t-1}}{\Pi_t}\right] \tag{29}$$

This wage rule has three key properties, similar to the wage rule in the main text. First, there is a positive relationship between wages and labor productivity, both in the aggregate component of labor productivity z_t and the idiosyncratic component of labor productivity x_{it}. Second, there is downward nominal wage rigidity. We assume that inflation is exogenous. Third, we fix wages over the course of the job, in order to isolate wages for newly hired workers. Wages do not change for continuing workers who have already been hired

E.1.3 Firm Problem

Our second departure from the standard DMP model is a model of the firm, similar to Elsby and Michaels (2013) and Michaillat (2012).

Firms produce output in order to maximize the present discount value of profits, taking prices as given. Each firm has a production function $y_{it} = z_t x_{it} n_{it}^{\alpha}$, with $\alpha \in (0,1)$. So, firms produce output with decreasing returns to scale, with aggregate labor productivity z_t and idiosyncratic labor productivity x_{it}. Since x_{it} is idiosyncratic, we have $\int_0^1 x_{it} di = 1$ for all t.

Firm profit is output, after deducting wages paid to new hires, vacancy posting costs for new hires, and wages paid to continuing workers, which we denote by C_{it}. So, firm i's profit in period t is

$$\pi_{it} = z_t x_{it} n_{it}^{\alpha} - w_{it} h_{it} - C_{it} - \frac{c}{q(\theta_t)} h_{it}. \tag{30}$$

Payroll paid to continuing workers satisfies

$$C_{it} = (1-s)\left(C_{i,t-1} + w_{i,t-1} h_{i,t-1}\right). \tag{31}$$

C_{it} is the sum of payroll paid to the previous group of continuing workers, and payroll paid to new hires from the previous period, after accounting for separations.

Hires and employment are linked by

$$n_{it} \leq (1-s)\, n_{i,t-1} + h_{it}. \tag{32}$$

Current employment is the previous stock of employment and current new hires, after accounting for separations. The firm problem is

$$\max_{\{h_{it}, n_{it}\}_{t=0}^{\infty}} E_0 \sum_{t-0}^{\infty} \beta^j \pi_{it} \tag{33}$$

subject to the wage rule (29), the equations of motion for continuing payroll (31), employment (32), and the non-negativity constraint on hiring h_{it}, while taking as given all aggregate variables and idiosyncratic labor productivity. h_{it} and n_{it} must be measurable with respect to time t information. For tractability in what follows, we assume that the solution to the firm's problem is always interior, and the non-negativity constraint on hiring never binds.

E.1.4 Equilibrium

In equilibrium, firms take as given:

1. A distribution of initial employment $n_{i,-1}$, idiosyncratic labor productivity $x_{i,-1}$, wages for new hires $w_{i,-1}$, and payroll paid to continuing workers $C_{i,-1}$
2. A process for aggregate and idiosyncratic labor productivity $\{z_t\}_{t=0}^{\infty}$ and and $\{x_{it}\}_{i=0}^{\infty}$

Then an equilibrium is a collection of stochastic processes $\{n_{it}, h_{it}, w_{it}, \theta_t, u_t\}_{t=0}^{\infty}$ which satisfies:

1. The firm's problem (33) for all i
2. The equation of motion for aggregate unemployment (28)
3. The wage rule (29)

E.2 Result: Job-Level Wages are Key for Unemployment Fluctuations

We now arrive at the main result. We show analytically that changes in job level wages are particularly important for unemployment fluctuations. This result justifies our focus in the empirics on job level wages. This result complements our argument in the main text, which makes the same point with a different extension of the standard DMP model.

We will consider an economy in aggregate steady state, for tractability. Firm-level variables remain uncertain. In the context of a single-firm DMP model Shimer (2005) argues this static environment reasonably approximates a fully dynamic DMP model. The reason is that in textbook calibrations of the DMP model, transition dynamics are rapid.

By introducing firms, we will have a meaningful notion of job level wages. A job level wage change is the change in a wage paid by a firm to new hires, across vacancies in successive time periods. We define job level wage changes as

$$\int_i \frac{dw_{it}}{dz} di$$

which is analogous to our measurement exercise.

We now argue that changes in the job level wage are particularly important for unemployment fluctuations To summarize unemployment fluctuations, we focussing on $d\log u/d\log z$, the elasticity of unemployment with respect to labor productivity. This elasticity measures the sensitivity of employment to labor demand shocks. As such, the elasticity is the primary focus of many past analytical results on unemployment in DMP models, such as Ljungqvist and Sargent (2017). We show that job level wage changes are a key determinant of the elasticity.

Proposition 3. *To a first order, the elasticity of aggregate employment to aggregate labor productivity is*

$$\frac{d\log u}{d\log z} = -\tilde{A} + \tilde{B}\int_0^1 \frac{dw_{it}}{dz} di$$

for constants $\tilde{A}, \tilde{B} > 0$.

Proof. Available on request. □

This result closely resembles equation (16) from the main text. The elasticity of unemployment to labor demand depends on an average of job-level wage changes. So, as in the model from the main text, job-level wage changes are particularly important for unemployment fluctuations. But in this section, we have proved the result with a different model, similar to Elsby and Michaels (2013) and Michaillat (2012).