Supplemental Appendix

Jonathon Hazell and Bledi Taska

A Additional Figures

Figure 1: Employment-Population Ratio in United States during 2019-2021

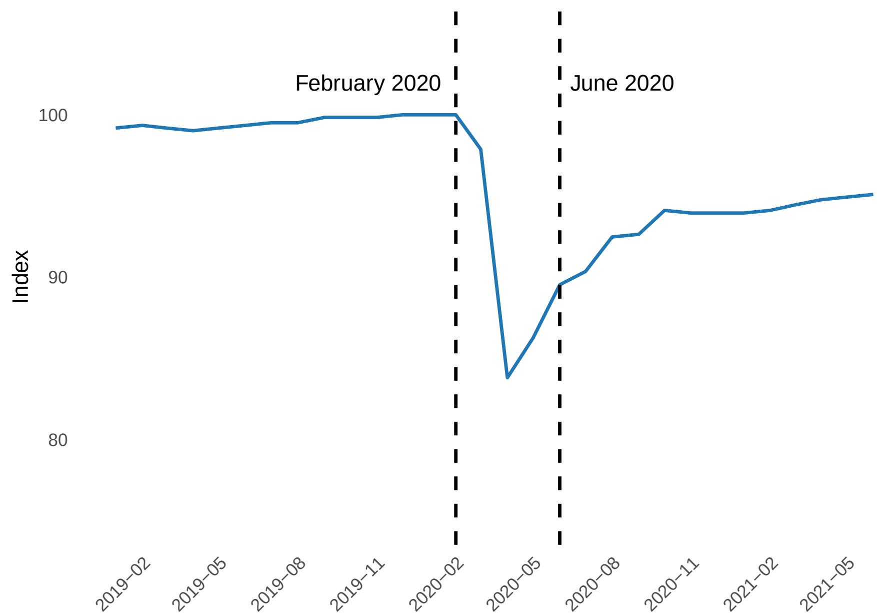

Notes: this graph plots the employent-population ratio, by month, from January 2019 to December 2021. The employment-population ratio is provided by the Bureau of Labor Statistics. The series is rescaled to have a value of 100 in February 2020.

Figure 2: Nominal Wages and Unemployment—pre vs. post June 2020

Notes: the graph plots binned wage growth of nominal posted wages, in percent, from Burning Glass; and binned state by quarter unemployment changes, in percentage points, from the Local Area Unemployment Statistics. In blue circles is data for 2010Q1-2020Q2, in red trianges is data for 2020Q3-2021Q1. 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 and add a non-parametric regression line, separately for each series.

Figure 3: Industry Employment Growth During the Pandemic Recession

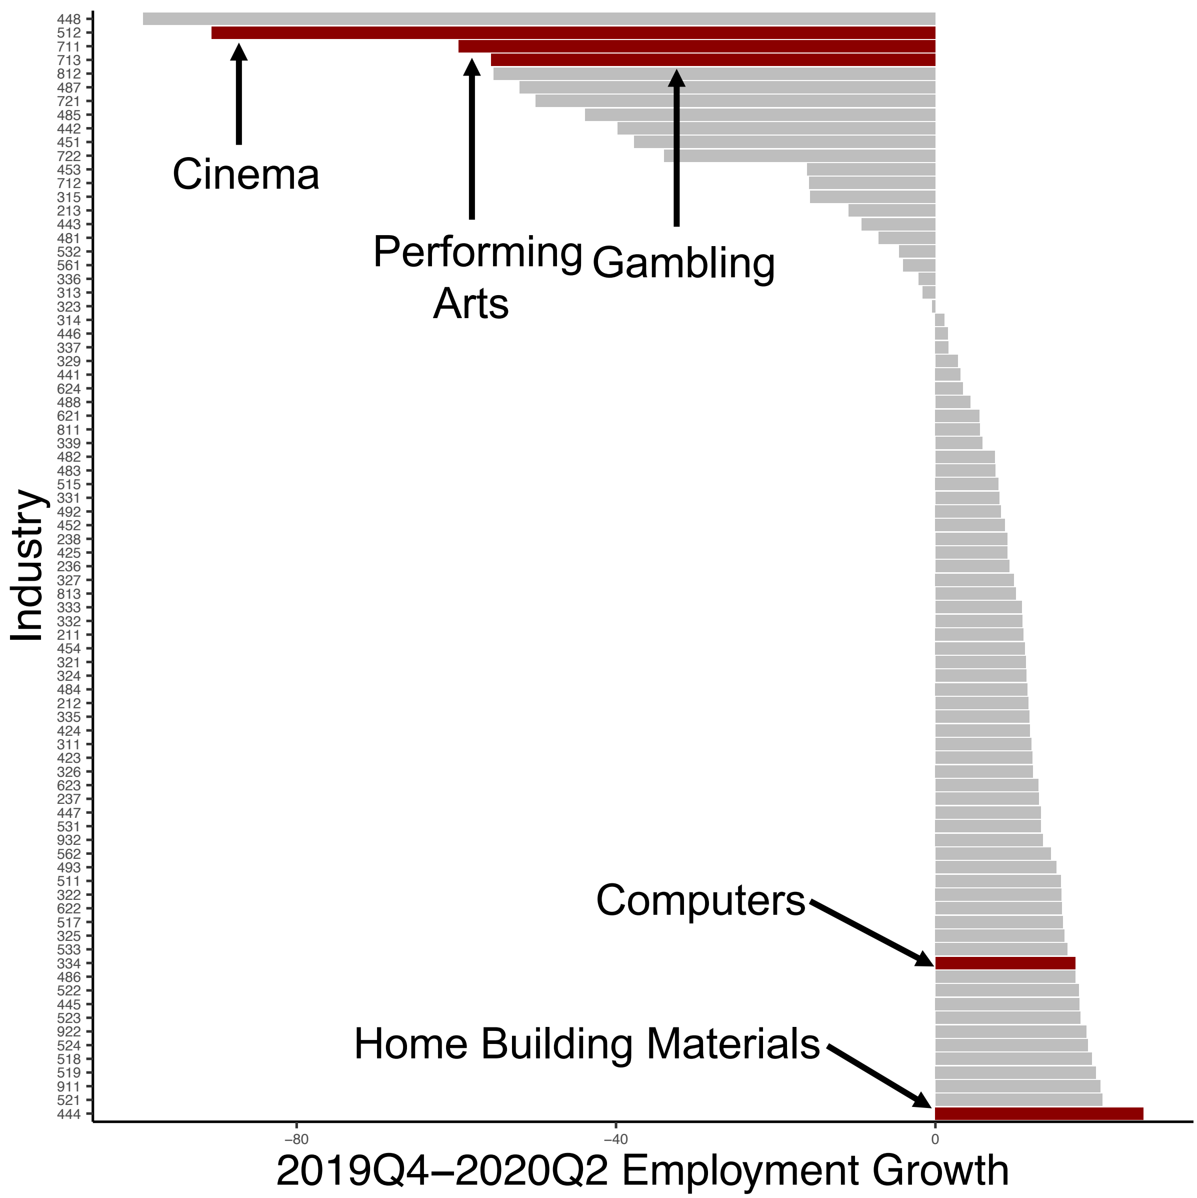

Notes: this graph plots industry employment growth, in percentage points, at the 3 digit NAICS level, from the Current Employment Statistics, between the last quarter of 2019 and the second quarter of 2020. Each observation is demeaned by national average employment growth over the period.

Figure 4: Brent Crude Oil Price ($/Barrel)

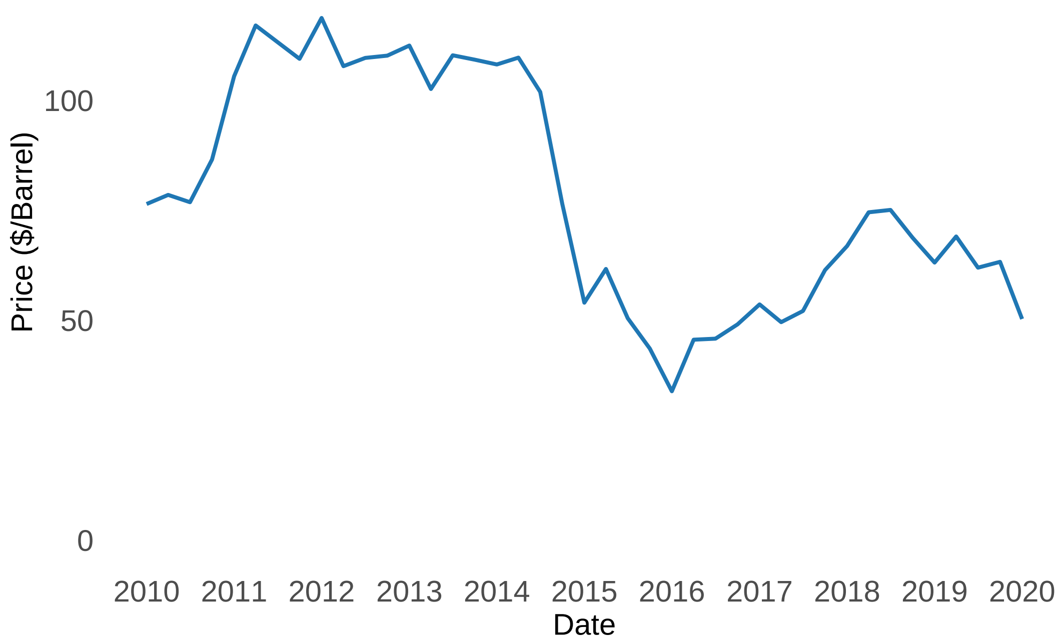

Notes: this graph plots the quarterly average of the Brent Crude oil price, for 2010-2020. The units are dollars per barrel of oil.

Figure 5: Nominal Wages and Industry Employment

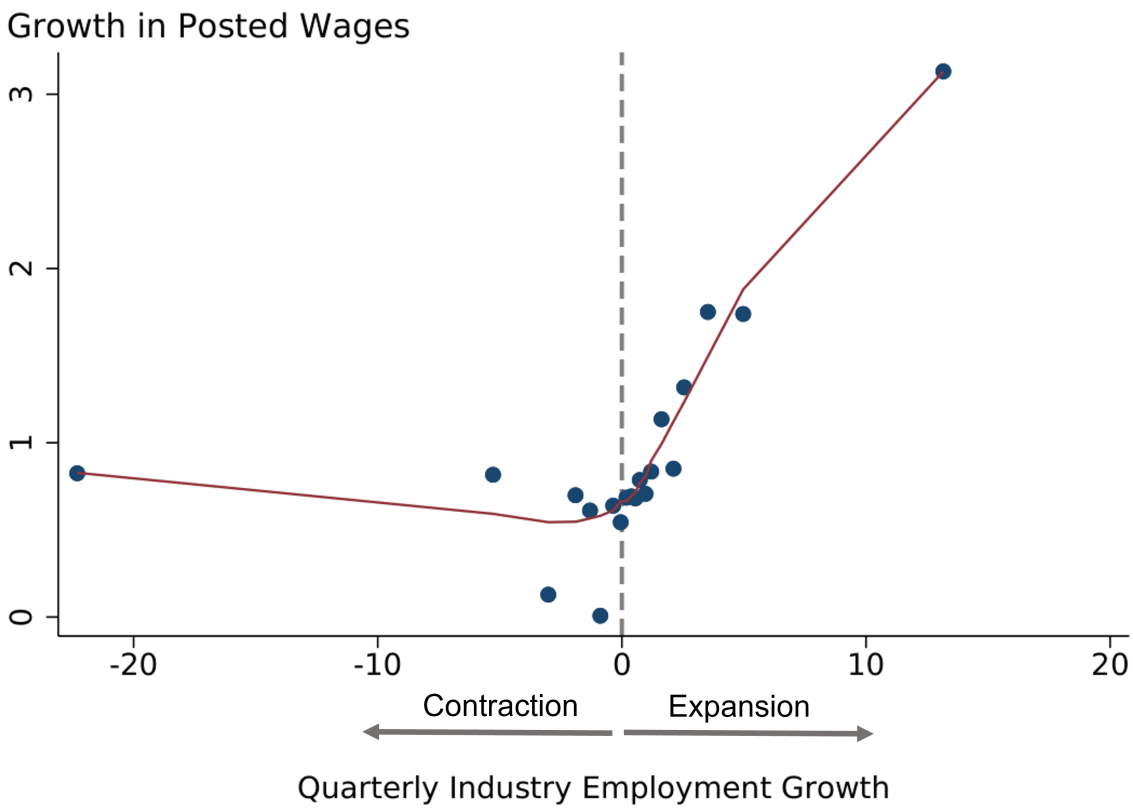

Notes: the graph plots binned wage growth of nominal posted wages, in percent, from Burning Glass; and binned 3 digit NAICS industry by quarter employment growth, in percentage points, from the Current Employment Statistics. The sample period is 2010Q1-2020Q2. To construct wage growth, we take the mean wage within each job and quarter, and then take log differences at the job level. We use 20 bins and add a non-parametric regression line.

Figure 6: State Level Coverage in Burning Glass

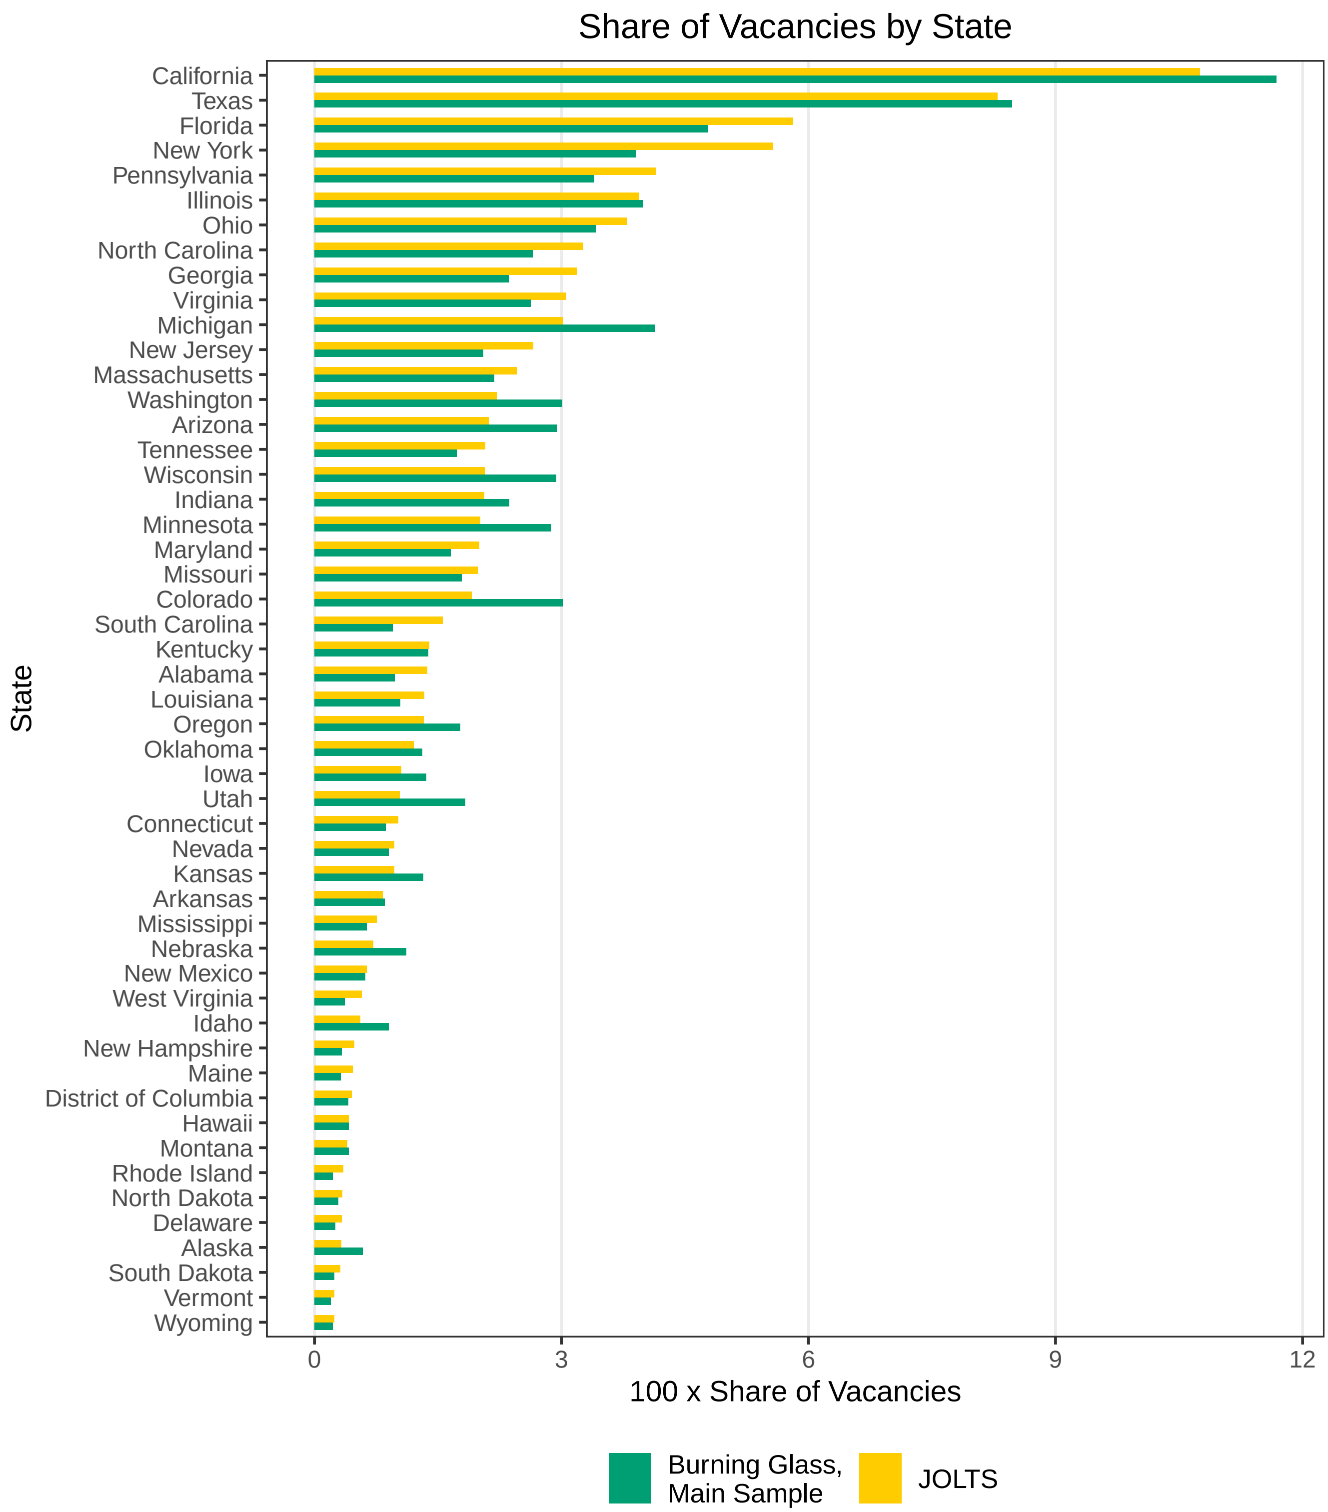

Notes: this graph compares the state level share of vacancies in Burning Glass and JOLTS. In green, is the share of vacancies in each state, in the Burning Glass main sample, for 2010-2020Q2. In yellow is the share of vacancies in each state, for 2010-2020, according to JOLTS.

Figure 7: Broad Occupation Coverage in Burning Glass

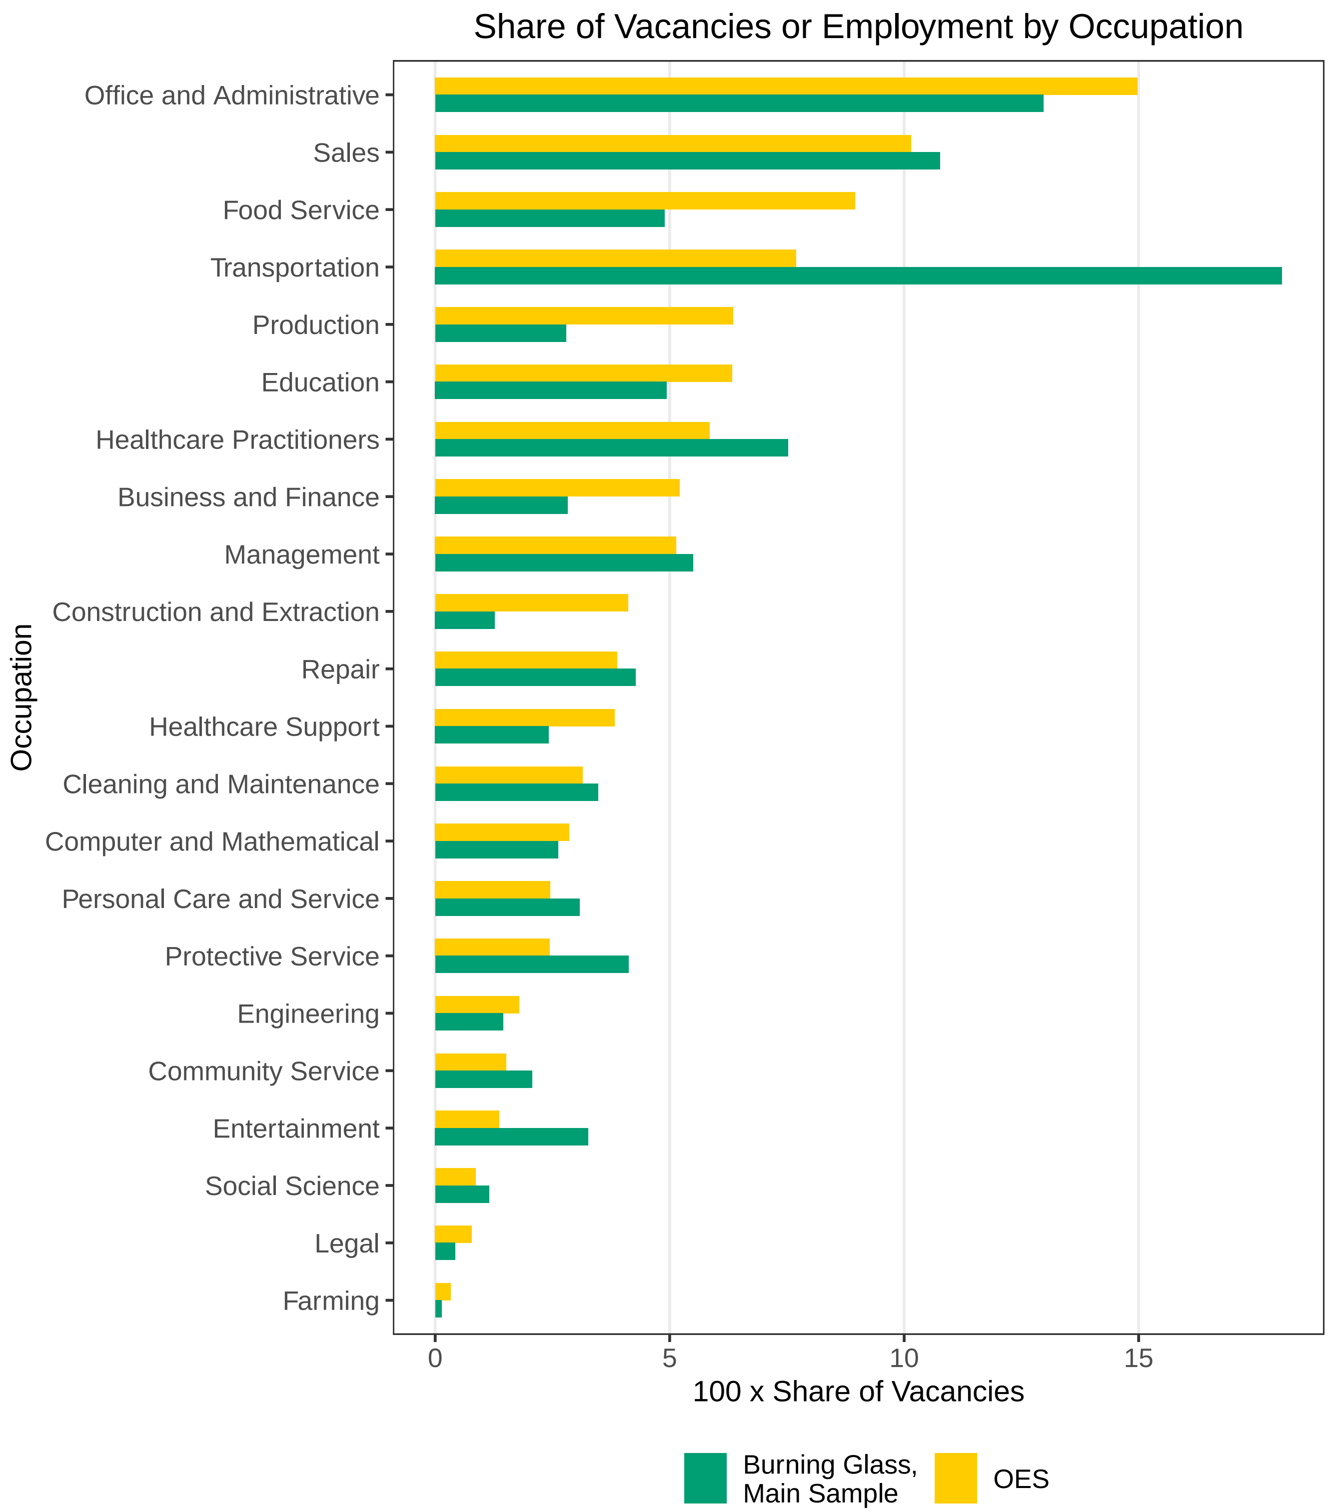

Notes: this graph compares the occupation level share of vacancies in Burning Glass versus employment in the Occupational Employment Statistics, at the 2 digit occupation level. In green, is the share of vacancies in each occupation, in the Burning Glass main sample, for 2010-2020Q2. In yellow is the share of employment in each occupation, for 2010-2020, according to the Occupational Employment Statistics.

Figure 8: State Level Coverage in Dun & Bradstreet and Burning Glass Merged Sample

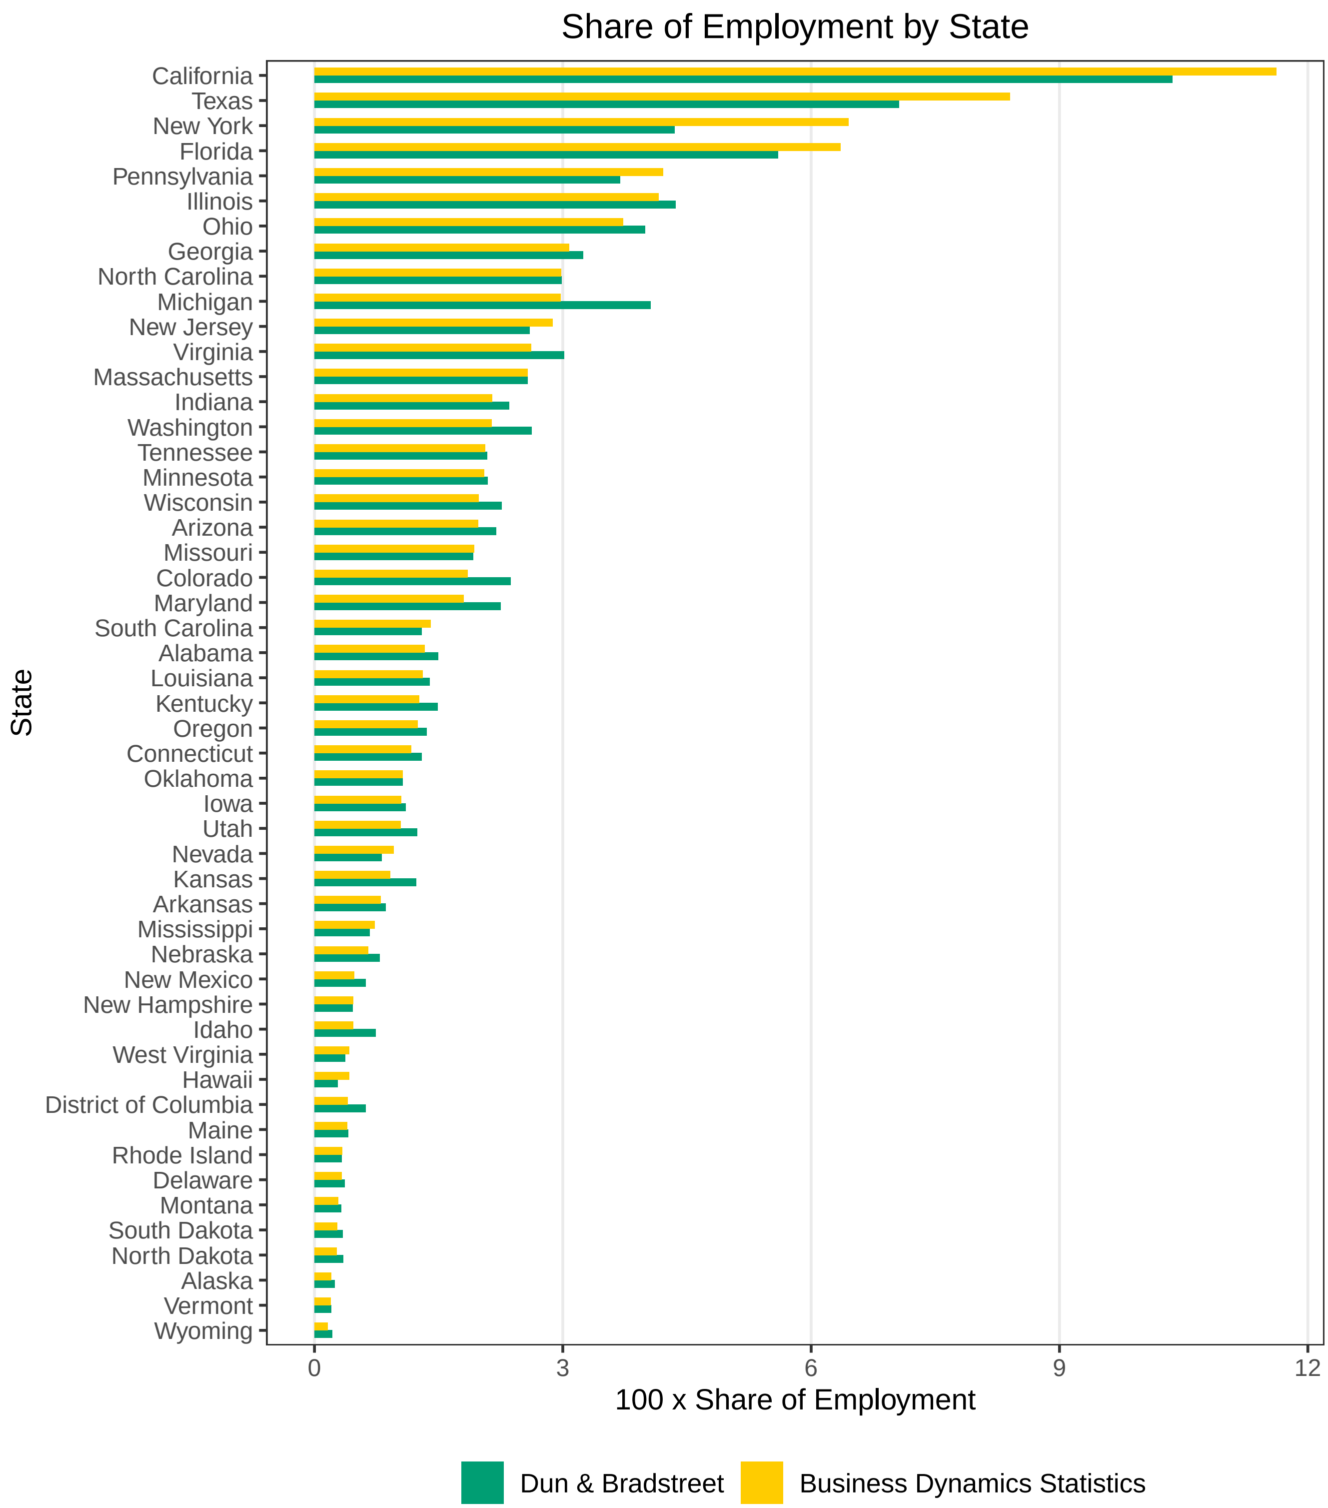

Notes: this graph compares the state level share of employment in Dun & Bradstreet and the Business Dynamics Statistics. In green, is the share of employment in each state, in the Dun & Bradstreet establishments that merge to the Burning Glass main sample, averaged over 2010-2020. In yellow is the share of employment in each state, according to the 2019 Business Dynamics Statistics.

Figure 9: Industry Coverage in Dun & Bradstreet and Burning Glass Merged Sample

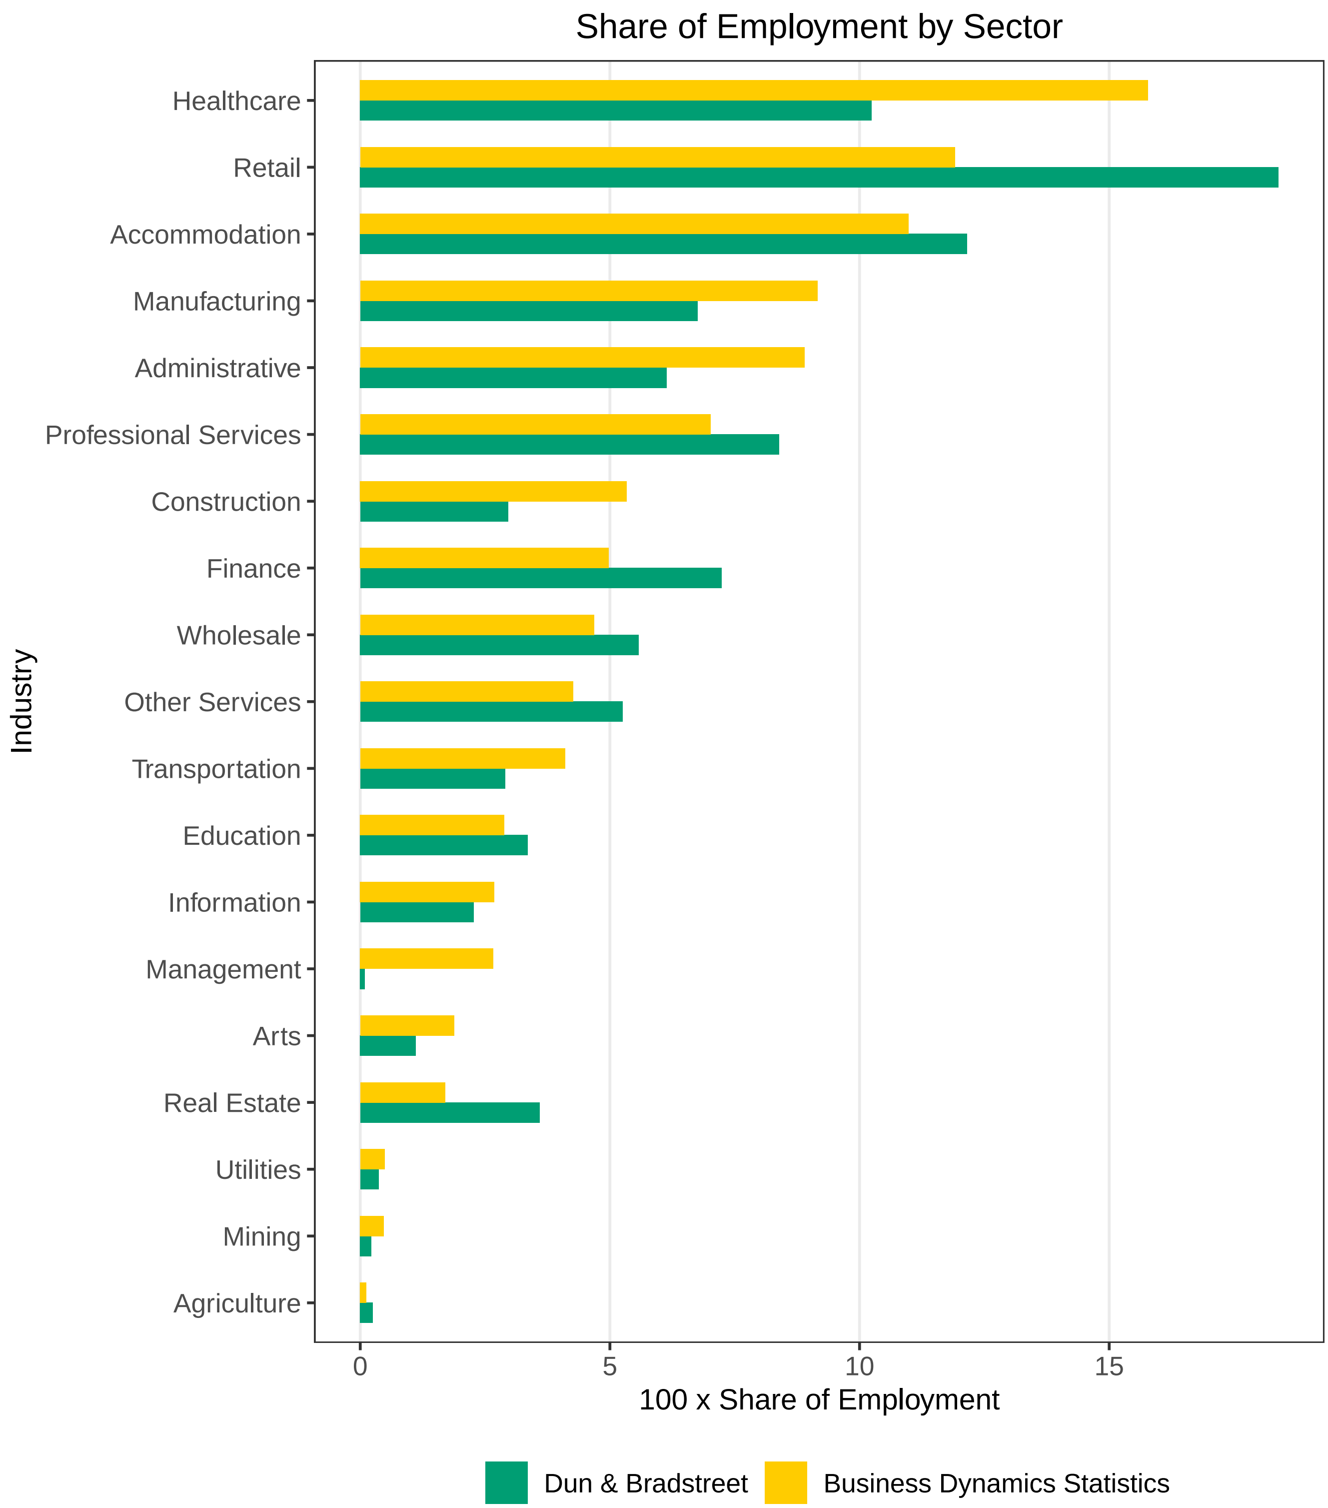

Notes: this graph compares the industry level share of employment in Dun & Bradstreet and the Business Dynamics Statistics, measured at the 2 digit industry level. In green, is the share of employment in each industry, for the Dun & Bradstreet establishments that merge to the Burning Glass main sample, averaged over 2010-2020. In yellow is the share of employment in each industry, according to the 2019 Business Dynamics Statistics.

Figure 10: Nominal Posted Wage Growth at the Job Level and Unemployment Changes-Heterogeneity by the Gap between Postings

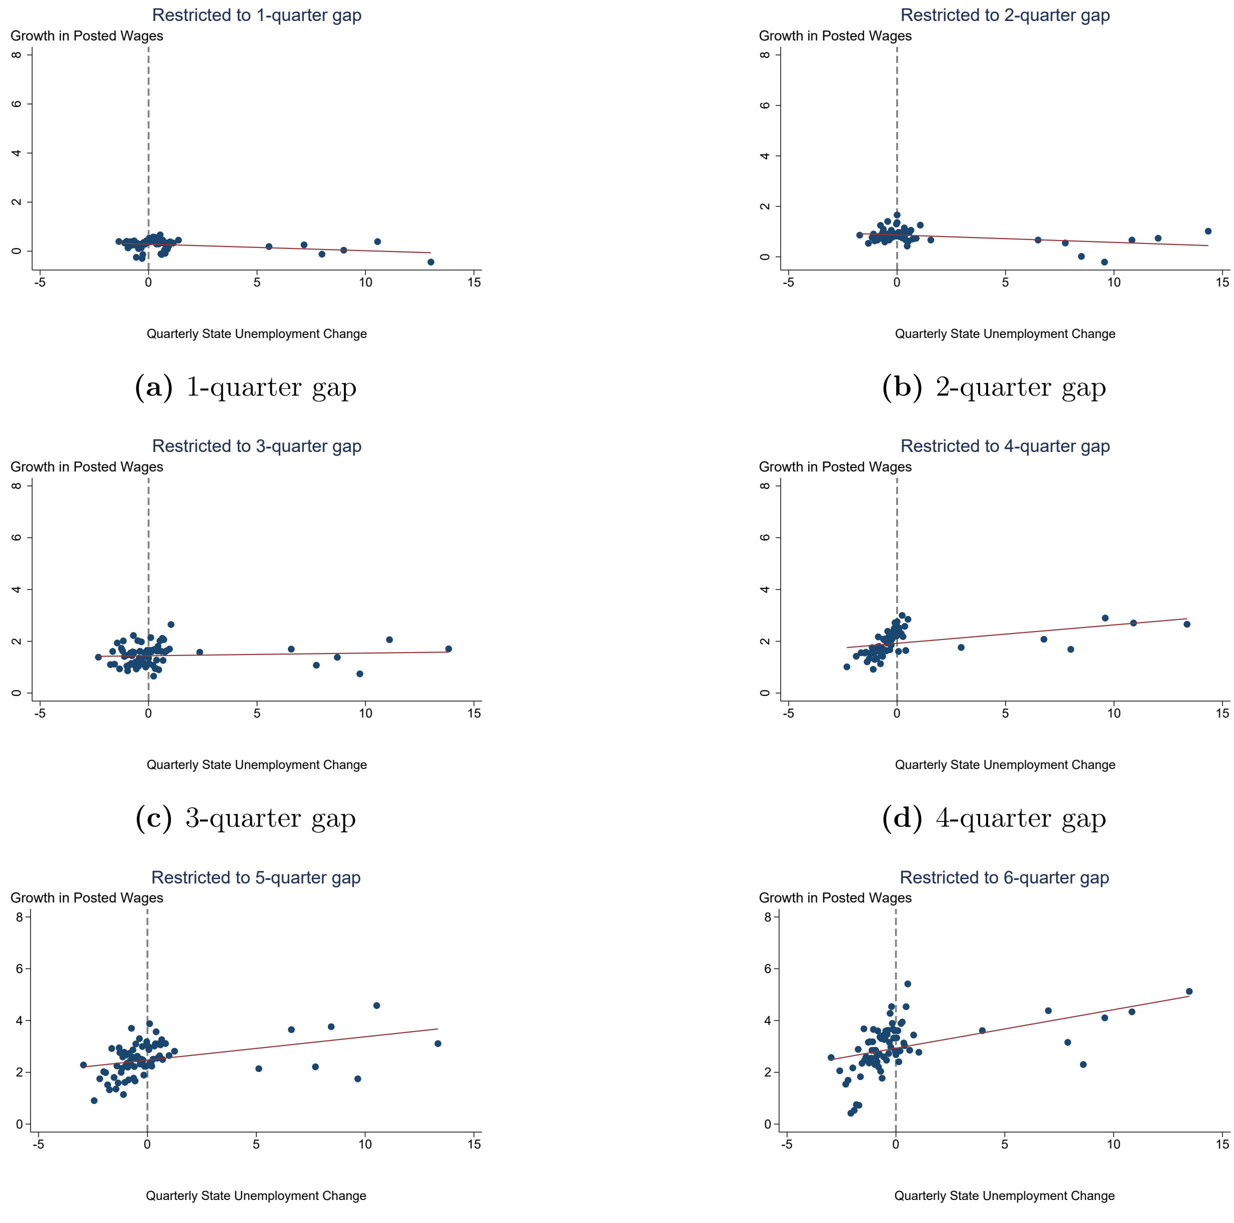

(a) 1-quarter gap

(b) 2-quarter gap

(c) 3-quarter gap

(d) 4-quarter gap

(e) 5-quarter gap

(f) 6-quarter gap

Note: the graph plots wage growth of nominal posted wages, in percent, from Burning Glass; and state by quarter unemployment changes, in percentage points, from the Local Area Unemployment Statistics. The sample period is 2010Q1-2020Q2. To construct wage growth, we take the mean wage within each job and quarter, and then take log differences at the job level. We collect wage growth and unemployment changes into 100 bins, and add a non-parametric regression line. Each panel corresponds to a specific temporal gap between job postings. Panel (a) restricts the sample to observations with a 1-quarter gap, Panel (b) to a 2-quarter gap, Panel (c) to a 3-quarter gap, Panel (d) to a 4-quarter gap, Panel (e) to a 5-quarter gap, and Panel (f) to a 6-quarter gap.

B Additional Tables

Table 1: Difference of State Wages in Burning Glass and CPS vs. Unemployment

	Difference between Quarterly State Wage Growth in Burning Glass and CPS			
	(1)	(2)	(3)	(4)
ΔU_{st}	0.12	0.13	0.06	0.06
	(0.38)	(0.38)	(0.38)	(0.38)
Observations	2,193	2,193	2,193	2,193
Time Fixed Effect	✓	✓	✓	✓
State Fixed Effect		✓		✓
Weights	QCEW	QCEW	CPS	CPS

Notes: this table regresses the difference in Burning Glass and Current Population Survey (CPS) wages on unemployment. The dependent variable is 100 times the growth of quarterly state wages in Burning Glass, minus 100 times the growth of quarterly state wages in the CPS. We measure quarterly state wages in Burning Glass and the CPS in the same way as Table 2, Panel A. The regressor is the change in state-quarter unemployment, from the 2010-2020 Local Area Unemployment Statistics, in percentage points. Standard errors are in parentheses, clustered by state. Columns (1) and (3) include time fixed effects, columns (2) and (4) include both time and state fixed effects. Column (1) and (2) are weighted by mean state employment during 2010-2020 from the Quarterly Census of Employment and Wages. Columns (3) and (4) are weighted by the quarterly size of the state, measured from the CPS.

Table 2: Variance Decomposition of Log Wages with Job Fixed Effects

Specification	Share of Variance Explained by Job FEs
Job FEs only	0.99
Job, Time FEs	0.98
Job, State-Time FEs	0.98
Job, State-Time FEs, at least 5 obs per job	0.94
Job FEs only, trimmed wages	0.98
Job title FEs	0.82

Notes: this table regresses log wages, for each job and quarter, on job fixed effects—recall that a job is a job title by establishment. Column (1) reports the specification, and column (2) reports the share of the total variance in wages that is explained by job fixed effects. Row 1 is the baseline specification. We add in time fixed effects in row 2, state-time fixed effects in row 3, and restrict to at least five observations per job in row 4 (recall that the main regression sample has at least two observations per job). Row 5 repeats row 1 but trims the most extreme 5% of wages within each year, 6 digit occupation, pay frequency and salary type. Row 6 repeats Row 1, but uses a different definition of a job. Specifically, we regress log wages only on job title fixed effects, instead of job title by establishment fixed effects. The regression is for the main Burning Glass sample, over 2010-2020Q2.

Table 3: Probability of Posted Wage Change By Gap Between Posting

Gap Between Postings (Quarters)	Quarterly Prob. Change	Quarterly Prob. Increase	Quarterly Prob. Decrease	Share of Observations (pp)
1	0.22	0.14	0.08	57
2	0.18	0.12	0.05	18
3	0.17	0.13	0.03	8
4	0.15	0.12	0.02	6
5	0.16	0.12	0.02	3
6	0.15	0.11	0.02	2
7	0.14	0.10	0.02	1
8	0.14	0.10	0.01	1

Notes: we study the main sample of Burning Glass data. We estimate the probability of job-level wage change (column 2), increase (column 3) and decrease (column 4) in the same way as table 3. Column 1 lists the number of quarters between the vacancy postings. In row (1), we estimate the probability only for vacancies that post in consecutive quarters (i.e. with a one quarter gap). In Row (2) we estimate the probability for vacancies that post with a two quarter gap, and so on. The final column reports the share of vacancies that post, given a number of quarters between postings.

Table 4: Quarterly Probability of Posted Wage Change—Heterogeneity by Length of Vacancy

	Prob. Change	Duration of Unchanged Wages	Prob. Decrease	Prob. Increase	Observations
	(1)	(2)	(3)	(4)	(5)
Above Median Vacancy Length	0.23	3.74	0.04	0.14	225,162
Below Median Vacancy Length	0.23	3.88	0.04	0.14	222,672

Notes: we study the main sample of Burning Glass data. We estimate the probability of job-level wage change (column 1), duration of unchanged wages (column 2) and probability of increase and decrease (columns 3 and 4) in the same way as table 3. In row (1), we estimate the probabilities for vacancies that post for an above median length of time. In row (2) we estimate the probabilities for vacancies that post for a below median length of time. The sample is vacancies from the main sample with information on the length of posting time.

Table 5: First Stage Regressions for Shift Share Instrument

	Δunemployment$_{ist}$			Δunemployment$_{ist} \times I\,(\Delta$unemployment$_{ist} < 0)$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log\,(\text{shift share}_{ist})$	-0.21	-0.33	-0.33	-0.30	-0.32	-0.32
	(0.02)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
$\Delta \log\,(\text{shift share}_{ist}) \times$	-0.54	-0.27	-0.26	0.30	0.31	0.31
$I\,(\Delta \log\,(\text{shift share}_{ist}) > 0)$	(0.03)	(0.05)	(0.05)	(0.03)	(0.02)	(0.02)
Observations	1,789,042	1,789,042	1,789,042	1,789,042	1,789,042	1,789,042
Time Fixed Effect		✓	✓		✓	✓
State Fixed Effect			✓			✓

Notes: this table presents first stage regressions of the shift share instrument. In Columns (1)-(3) the outcome variable is the change in quarter by state unemployment, from the 2010-2020Q2 Local Area Unemployment Statistics. In Columns (4)-(6) the outcome variable is the change in unemployment interacted with an indicator for whether unemployment is declining. In all columns, the regressors are the quarterly change in the shift share instrument at the job level, and the change in the shift share interacted with an indicator for whether the shift share is increasing. In columns (1) and (4) there are no controls. Columns (2) and (5) have time fixed effects interacted with an indicator for whether the instrument is positive. Columns (3) and (6) add state fixed effects. Standard errors are in parentheses, clustered by state.

Table 6: Nominal Posted Wages and Unemployment—Heterogeneity by Posting Time

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
Vacancy posting time:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ΔU_{st}	-0.11	-0.09	-0.02	0.01
	(0.07)	(0.06)	(0.03)	(0.04)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-2.55	-2.61	-2.98	-2.78
	(0.25)	(0.24)	(0.24)	(0.18)
Observations	112,642	114,682	109,081	110,981
Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity by the length of vacancy posting. We estimate the baseline regression, that is, column (2) of Table 4. We estimate wage rigidity separately for vacancies according to the length of time that vacancies were posted, split into quartiles. The sample is vacancies from the main sample with information on the length of time for which vacancies were posted.

Table 7: Heterogeneity in Wage Cyclicality by Occupation

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$				
	Management	Services	Sales	Construction	Production
ΔU_{st}	-0.10	0.02	-0.04	-0.07	-0.05
	(0.03)	(0.03)	(0.04)	(0.06)	(0.05)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.30	-2.03	-1.63	-1.80	-2.19
	(0.14)	(0.18)	(0.11)	(0.16)	(0.15)
Observations	549,934	312,356	399,422	99,114	356,921
Time Fixed Effect	✓	✓	✓	✓	✓
Previous Time Fixed Effect	✓	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity by occupation. We estimate the baseline regression, that is, column (2) of Table 4. We estimate wage rigidity separately for five broad occupations according to the SOC classification. Standard errors are in parentheses, clustered by state.

Table 8: Heterogeneity in Wage Cyclicality by Industry

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$				
Panel A:	Natural Resources	Construction	Manufacturing	Trade & Utilities	Information
ΔU_{st}	0.37	0.02	-0.08	-0.09	-0.17
	(0.06)	(0.10)	(0.05)	(0.09)	(0.16)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.86	-1.25	-1.67	-2.18	-1.45
	(0.22)	(0.25)	(0.10)	(0.15)	(0.20)
Observations	11,533	20,032	82,893	365,584	48,959
Panel B:	Finance	Business Services	Education & Health	Leisure	Other Services
ΔU_{st}	-0.03	-0.01	-0.08	0.04	-0.08
	(0.03)	(0.03)	(0.03)	(0.06)	(0.05)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.91	-1.72	-1.48	-2.16	-1.60
	(0.12)	(0.19)	(0.13)	(0.16)	(0.29)
Observations	123,932	139,969	456,709	112,697	27,808
Time Fixed Effect	✓	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity by industry. We estimate the baseline regression, that is, column (2) of Table 4. We estimate wage rigidity separately for ten broad industries, at the 2 digit NAICS level. Standard errors are in parentheses, clustered by state.

Table 9: Nominal Posted Wages and Unemployment at the Establishment Level

	Nominal Wage Growth at Establishment Level, $\Delta \log w_{jst}$			
	(1)	(2)	(3)	(4)
ΔU_{st}	-0.09	-0.14	-0.13	-0.81
	(0.01)	(0.06)	(0.06)	(0.11)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.45	-1.76	-1.79	
	(0.16)	(0.19)	(0.21)	
Observations	2,120,212	2,120,212	2,120,212	2,120,212
Time Fixed Effect		✓	✓	✓
State Fixed Effect			✓	

Notes: this table presents estimates regressing nominal establishment-level wages on unemployment. The dependent variable is quarterly percent growth in nominal posted wages, from the Burning Glass main sample. Wage growth is trimmed at the 1st and 99th percentiles. The regressors are the change in state-quarter unemployment, and the change interacted with an indicator for whether unemployment is decreasing, from the 2010-2020 LAUS, in percentage points. Standard errors are in parentheses, clustered by state. Column (1) presents estimates without controls. Column (2) adds in time fixed effects, interacted with an indicator variable for whether unemployment is decreasing. Column (3) adds in state fixed effects. Column (4) presents estimates without asymmetries, by only including unemployment changes and time fixed effects as regressors.

Table 10: Nominal Posted Wages and Unemployment—Heterogeneity by Posting Frequency

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
Number of vacancy postings:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ΔU_{st}	-0.05	-0.05	-0.02	-0.07
	(0.04)	(0.02)	(0.02)	(0.03)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.71	-1.75	-1.72	-1.37
	(0.12)	(0.14)	(0.10)	(0.11)
Observations	728,464	396,696	365,447	298,434
Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity by the frequency of vacancy posting. We estimate the baseline regression, that is, column (2) of Table 4. We estimate wage rigidity separately for vacancies according to the number of times that vacancies were posted, split into four groups. Note that although we use the word quartile in the table, we use the jobs that posted two times in first column and jobs that posted three times in the second column. In the last two columns, we split the rest of the sample evenly into two groups.

Table 11: Heterogeneity in Wage Cyclicality by Source of Vacancy

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
	Company Website	Job Board	Government	Education
ΔU_{st}	-0.16	-0.17	1.17	-0.30
	(0.20)	(0.14)	(0.93)	(0.50)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.63	-1.35	-2.20	-1.16
	(0.24)	(0.11)	(0.77)	(0.56)
Observations	120,277	325,838	51,985	51,070
Time Fixed Effect	✓	✓	✓	✓
Previous Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity by source of vacancy. We estimate the baseline regression, that is, column (2) of Table 4. We estimate wage rigidity separately for vacancies from different sources, as listed in the table. The sample is vacancies from the main sample with information on the source of the vacancy. Standard errors are in parentheses, clustered by state.

Table 12: Regression of Nominal Posted Wage Growth on Industry Employment Growth

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jit}$			
	(1)	(2)	(3)	(4)
$\Delta \log(\text{employment}_{it})$	0.00	-0.01	-0.01	0.01
	(0.00)	(0.00)	(0.00)	(0.00)
$\Delta \log(\text{employment}_{it}) \times$	0.05	0.06	0.05	
$I(\Delta \log(\text{employment}_{it}) > 0)$	(0.02)	(0.01)	(0.01)	
Observations	934,502	934,502	934,502	934,502
Time Fixed Effect		✓	✓	✓
Industry Fixed Effect			✓	

Notes: this table presents estimates regressing nominal job-level wages on industry employment, at the 3 digit NAICS level. The dependent variable is quarterly percent growth in nominal posted wages, from the Burning Glass main sample. Wage growth is trimmed at the 1st and 99th percentiles. The regressors are the growth in industry by quarter employment, and the change interacted with an indicator for whether employment is increasing, from the 2010-2020 Current Employment Statistics, in percentage points. Standard errors are in parentheses, clustered by industry. Column (1) presents estimates without controls. Column (2) adds in time fixed effects, interacted with an indicator variable for whether employment is increasing. Column (3) adds in industry fixed effects. Column (4) presents estimates without asymmetries, by only including employment changes and time fixed effects as regressors.

Table 13: Heterogeneity in Wage Cyclicality by Establishment Size and Occupation Wage

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
Position in size or wage distribution:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	Panel A: Heterogeneity by Establishment Size			
ΔU_{st}	0.01	-0.03	-0.17	-0.08
	(0.02)	(0.02)	(0.09)	(0.04)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.71	-1.70	-1.69	-1.55
	(0.09)	(0.09)	(0.14)	(0.17)
Observations	493,623	409,625	439,473	446,320
	Panel B: Heterogeneity by Occupation Wage			
ΔU_{st}	-0.02	-0.09	-0.03	-0.06
	(0.05)	(0.04)	(0.02)	(0.03)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-2.04	-1.70	-1.75	-1.27
	(0.17)	(0.11)	(0.13)	(0.13)
Observations	429,883	393,026	406,396	402,585
Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity. We estimate the baseline regression, that is, column (2) of Table 4. In Panel A, we estimate wage rigidity separately for vacancies in the first through fourth quartiles of the establishment size distribution. Establishment size is the total number of vacancies posted by an establishment, in the main sample, during 2010-2020. In Panel B, we estimate wage rigidity separately for vacancies in the first through fourth quartiles of the wage distribution. Wages are the median wage, within the 6 digit occupation, measured from the 2014-16 Occupational Employment Statistics. Standard errors are in parentheses, clustered by state.

Table 14: Change in Share of Vacancies with Wages at Establishment by Occupation Level

	Change in Share of Vacancies with Posted Wages			
	(1)	(2)	(3)	(4)
ΔU_{st}	-0.00	0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	0.00	0.00	0.00	
	(0.00)	(0.00)	(0.00)	
Observations	2,975,670	2,975,670	2,975,670	2,975,670
Time Fixed Effect		✓	✓	✓
State Fixed Effect			✓	

Notes: this table presents estimates regressing the change in the share of vacancies with wages on the change in unemployment. The dependent variable is quarterly change in the share of vacancies that post wages, within each establishment and 6 digit occupation. The regressors are the change in state-quarter unemployment, and the change interacted with an indicator for whether unemployment is decreasing, from the 2010-2020 LAUS, in percentage points. Standard errors are in parentheses, clustered by state. Column (1) presents estimates without controls. Column (2) adds in time fixed effects, interacted with an indicator variable for whether unemployment is decreasing. Column (3) adds in state fixed effects. Column (4) presents estimates without asymmetries, by only including unemployment changes and time fixed effects as regressors. The sample is all establishment by occupation observations that post at least one vacancy with a wage in the lagged period.

Table 15: Heterogeneity in Wage Cyclicality by Degree of Wage Bargaining or Posting

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
Degree of bargaining or posting in occupation:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	Panel A: Heterogeneity by Degree of Wage Bargaining			
ΔU_{st}	-0.09	-0.03	-0.04	-0.02
	(0.03)	(0.04)	(0.05)	(0.04)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.32	-1.69	-2.16	-1.91
	(0.14)	(0.11)	(0.15)	(0.13)
Observations	530,381	517,557	362,626	307,183
	Panel B: Heterogeneity by Degree of Wage Posting			
ΔU_{st}	-0.04	-0.05	-0.02	-0.08
	(0.05)	(0.05)	(0.03)	(0.03)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.96	-1.62	-1.73	-1.47
	(0.15)	(0.11)	(0.12)	(0.14)
Observations	458,915	424,497	413,320	421,013
Interacted Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity. We estimate the baseline regression, that is, column (2) of Table 4. In Panel A, we estimate wage rigidity separately for vacancies in the first through fourth quartiles of the share of workers within an occupation that engage in wage bargaining. In Panel B, we estimate wage rigidity separately for vacancies in the first through fourth quartiles of the share of workers within an occupation that receive a job with a posted wage. The share of workers within an occupation that either receive a bargained or posted wage is provided by Hall and Krueger (2012). Standard errors are in parentheses, clustered by state.

Table 16: Regression of State Share of High Wage Jobs on Unemployment

	Change in State Share of High Wage Jobs			
	(1)	(2)	(3)	(4)
ΔU_{st}	0.11	0.09	-0.05	-0.07
	(0.18)	(0.18)	(0.17)	(0.17)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$			0.73	0.84
			(0.86)	(0.96)
Observations	2,040	2,040	2,035	2,035
Time Fixed Effect	✓	✓	✓	✓
State Fixed Effect		✓		✓

Notes: In all columns, the dependent variable is the growth in the share of high wage vacancies in the state, in percentage points. High wage vacancies have wages above the national median, within the pay frequency and salary type. In columns (1) and (2), the independent variable is state unemployment changes, in percentage points, from the 2010-2020 LAUS. Column (2) also controls for state trends. Columns (3) and (4) repeat columns (1) and (2), but include an additional regressor, the interaction of unemployment changes with an indicator for whether unemployment is falling. We also control for the interaction of this indicator with a time fixed effect. We weight by mean state employment over 2010-2020 from the QCEW. Standard errors are in parentheses, clustered by state.

Table 17: Regression of Average Nominal Wage Growth on Unemployment Changes

	Change in State Average of		Change in National Average of Wage for New Hires in	
	Posted Wages in Burning Glass	Wage for New Hires in Current Population Survey	Current Population Survey	National Longitudinal Survey of Youth
	(1)	(2)	(3)	(4)
ΔU_{st}	-0.29	-0.10	3.77	-1.78
	(0.16)	(0.59)	(3.47)	(3.17)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	0.40	0.38	-5.11	2.94
	(0.78)	(1.56)	(5.15)	(4.31)
Observations	2,085	1,627	83	89
Time Fixed Effect	✓	✓	✓	✓
Occupation FE		✓		✓
Industry FE		✓		✓
Demographics		✓	✓	✓

Notes: In column (1), the dependent variable is the percent growth in average state wages, from Burning Glass, measured in the same way as Table 2. The independent variable is state unemployment changes, in percentage points, from the 2010-2020 LAUS, with time fixed effects; as well as the interaction of unemployment changes with an indicator for whether unemployment is falling, and the interaction of this indicator with a time fixed effect. Column (2) repeats Column (1) but replaces the dependent variable with the percentage growth in average state wages, from the Current Population Survey, first we residualize wages against 3 digit occupation and industry fixed effects, and demographic controls (gender, race, marital status, education and a fourth order polynomial in experience) and then take the mean. Columns (1) and (2) are weighted mean state employment over 2010-2020 from the QCEW, and standard errors are clustered by state. In column (3), the dependent variable is the percent growth in average wages for new hires, from the Current Population Survey, for 1984-2007, as constructed by Haefke et al (2013). The independent variable is national unemployment changes, in percentage points; as well as the interaction of unemployment changes with an indicator for whether unemployment is falling. Column (4) repeats Column (3), but replaces the outcome with the measure of the wage for new hires constructed by Basu and House (2016) from the National Longitudinal Survey of Youth; this measure controls for occupation, industry and demographic characteristics, as well as the cumulative tightness measure of Hagedorn and Manovskii (2013). Columns (3) and (4) use robust standard errors.

Table 18: Correcting Selection Bias—Robustness

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$					
	Control for employment growth	Trim probabilities	Lagged wage in selection equation	Including entrants	Establishment Level	Annual
ΔU_{st}	-0.05	-0.01	-0.06	-0.00	-0.02	-0.12
	(0.01)	(0.02)	(0.02)	(0.02)	(0.05)	(0.08)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.78	-1.45	-1.66	-1.48	-0.87	-1.38
	(0.10)	(0.07)	(0.11)	(0.07)	(0.10)	(0.14)
Observations	1,450,049	1,448,998	610,334	1,458,049	798,460	686,647
Time Fixed Effect	✓	✓	✓	✓	✓	✓
Previous Time Fixed Effect	✓	✓	✓	✓	✓	✓

Notes: this table reports robustness of the main selection corrected regression specification, column (2) of Table 8. In the first column, we remove the Heckman selection terms, and instead control for employment growth between the start of the year and the start of the previous year, in the establishment. In the second column, we repeat the baseline Heckman correction, but drop observations with a predicted probability of selection outsize the unit interval, following the recommendation of Das et al (2003). In column (3), we re-estimate the baseline Heckman correction but add the lagged wage to the variables in the selection equation. In column (4) we also include entrants. In columns (5) and (6), we aggregate the selection equations and selection-corrected regressions to the establishment and annual level, respectively. Standard errors are in parentheses, clustered by state.

Table 19: Heterogeneity in Wage Cyclicality by Turnover

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
Ratio of vacancies to employment:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
ΔU_{st}	-0.08	-0.04	-0.06	-0.01
	(0.02)	(0.02)	(0.08)	(0.02)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.65	-1.70	-1.55	-1.80
	(0.07)	(0.13)	(0.15)	(0.16)
Observations	434,159	473,787	434,974	374,827
Time Fixed Effect	✓	✓	✓	✓
Previous Time Fixed Effect	✓	✓	✓	✓

Notes: this table reports estimates of downward wage rigidity with heterogeneity. We estimate the baseline regression, that is, column (2) of Table 4. We calculate the ratio of vacancies to employment within each occupation. Vacancies are the number of observations in each 6 digit occupation, for the main sample, and employment is measured at the 6 digit level from the 2014-16 Occupational Employment Statistics. W estimate wage rigidity separately for observations in the first through fourth quartiles of the distribution of the ratio of vacancies to employment. Standard errors are in parentheses, clustered by state.

Table 20: Detailed Occupations in Burning Glass vs. Occupational Employment Statistics

	6 Digit SOC Occupation Shares in Burning Glass		
	(1)	(2)	(3)
Occupation Shares in OES	0.84	0.86	0.92
	(0.05)	(0.05)	(0.06)
Observations	666	666	666
2 Digit Occupation Fixed Effect		✓	
3 Digit Occupation Fixed Effect			✓

Notes: the outcome variable is the share of vacancies within each 6 digit occupation, for the main Burning Glass sample. The regressor is the share of employment within each 6 digit occupation, from the 2014-16 Occupational Employment Statistics. Column (1) has no controls, column (2) adds 2 digit occupation fixed effects, and column (3) adds 3 digit occupation fixed effects. Robust standard errors are in parentheses.

Table 21: County Employment in County Business Patterns vs. in Burning Glass

	Log County Vacancies in Burning Glass		
	(1)	(2)	(3)
Log County Employment	0.98	0.99	0.98
	(0.01)	(0.01)	(0.01)
Observations	3,013	3,013	3,013
Census Division Fixed Effect		✓	
State Fixed Effect			✓

Notes: the outcome variable is 100 times the log number of vacancies within each county, for the main Burning Glass sample. The regressor is 100 times log employment for each county, as measured by the 2016 County Business Patterns. Column (1) has no controls, column (2) adds fixed effects for the census division, and column (3) adds fixed effects for the state. Robust standard errors are in parentheses.

Table 22: Testing Whether Coverage of Dataset is Cyclical

	Change in Share of Coverage at the State by Quarter Level			
	(1)	(2)	(3)	(4)
	Panel A: Outcome is Δ (Total Burning Glass Vacancies$_{st}$/JOLTS Vacancies$_{st}$)			
ΔU_{st}	0.10	0.11	0.15	0.16
	(0.35)	(0.35)	(0.36)	(0.36)
Observations	2,040	2,040	2,040	2,040
	Panel B: Outcome is Δ (Burning Glass Vacancies in Main Sample$_{st}$/Total Burning Glass Vacancies$_{st}$)			
ΔU_{st}	0.02	0.02	0.02	0.02
	(0.01)	(0.01)	(0.01)	(0.01)
Observations	2,091	2,091	2,091	2,091
Time Fixed Effect	✓	✓	✓	✓
State Fixed Effect		✓		✓
Weights	Employment	Employment	Vacancies	Vacancies

Notes: this table studies whether the coverage of various datasets is cyclical. In Panel A, the outcome variable is the change in the share of total 2010-2020Q2 Burning Glass vacancies in each quarter and state, relative to seasonally adjusted total vacancies in each state and quarter from the seasonally adjusted 2010-2020Q2 Job Opening and Labor Turnover Survey, measured in percentage points. We apply the rescaling formula of Carnevale et al (2014) so that vacancies in JOLTS and Burning Glass are expressed in the same units. In Panel B, the outcome variable is the change in the share of Burning Glass vacancies in each quarter and state that are in the main regression sample, relative to total vacancies in Burning Glass in the state and quarter, measured in percentage points.The main regression sample is vacancies with job title information, that post a point wage, and post in multiple quarters. In all panels, the regressor is quarter by state unemployment changes, in percentage points, from the 2010-2020Q2 Local Area Unemployment Statistics. Columns (1) and (2) are weighted by average state employment for 2010-2020 from the Quarterly Census of Employment and Wages. Columns (3) and (4) are weighted by average state vacancies for 2010-2020 from the full sample of Burning Glass vacancies. Standard errors are in parentheses, clustered by state.

Table 23: Testing Whether Coverage of Dataset Varies by Urban Status

	Change in Share of Coverage at the State by Quarter Level, Heterogeneity by Urban Share			
	(1)	(2)	(3)	(4)
	Panel A: Outcome is Δ (Total Burning Glass Vacancies$_{st}$/JOLTS Vacancies$_{st}$)			
$\Delta U_{st} \times I$ (low urban share)	0.40	0.42	0.53	0.55
	(0.39)	(0.39)	(0.38)	(0.39)
$\Delta U_{st} \times I$ (high urban share)	0.08	0.09	0.12	0.14
	(0.30)	(0.30)	(0.31)	(0.31)
Observations	2,040	2,040	2,040	2,040
	Panel B: Outcome is Δ (Burning Glass Vacancies in Main Sample$_{st}$/Total Burning Glass Vacancies$_{st}$)			
$\Delta U_{st} \times I$ (low urban share)	0.02	0.02	0.02	0.02
	(0.01)	(0.01)	(0.01)	(0.01)
$\Delta U_{st} \times I$ (high urban share)	0.02	0.02	0.02	0.02
	(0.01)	(0.01)	(0.01)	(0.01)
Observations	2,091	2,091	2,091	2,091
Time Fixed Effect	✓	✓	✓	✓
State Fixed Effect		✓		✓
Weights	Employment	Employment	Vacancies	Vacancies

Notes: this table repeats the analysis of 22, while modifying the regressors. In this table, the regressors are the change in quarterly state unemployment, interacted with whether the state has an above or below weighted median share of the population in urban areas. We rank states by their share in urban areas, as measured by the Census Bureau in 2010. The median is weighted by average state employment over 2010-2020.

Table 24: Statistics for Dun & Bradstreet / Burning Glass Merged Sample

	Business Dynamics Statistics	Burning Glass / Dun & Bradstreet
Panel A: Establishment Age (Years)		
0-5	20	14
6-10	14	24
11+	66	62
Panel B: Establishment Employment		
1-19	24	6
20-499	54	56
500+	22	38
Panel C: Establishment Idiosyncratic Shocks (pp)		
Job creation rate	12	13
Job destruction rate	10	16
Reallocation rate	20	27

Notes: this table compares information from the 2019 Business Dynamics Statistics (BDS), to information on establishments in Dun & Bradstreet that merge to the main Burning Glass sample. In column (1) is information from the BDS, and in column (2) is the corresponding information from the Burning Glass main sample. Panel A reports the share of establishments in each age bucket, where the buckets are 0-5 years in operation, 6-10 years in operation, and at least 11 years in operation. Panel B reports the share of establishments in each size bucket, where the buckets are 1-19 employees, 20-499 employees, and at least 500 employees. Panel C reports measures of establishment idiosyncratic shocks. In column (1) we report the job creation rate, job destruction rate and reallocation rate,as reported by the BDS and originally defined by Davis, Haltiwanger, and Schuh (1998). In column (2) we report the median job creation, job destruction and reallocation rates from the Burning Glass main sample.

Table 25: Wage Growth and Unemployment Changes, Bad Control

	Nominal Wage Growth at the Job Level, $\Delta \log w_{jst}$			
	(1)	(2)	(3)	(4)
ΔU_{st}	0.05	0.01	-0.02	0.01
	(0.01)	(0.03)	(0.02)	(0.03)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.40	0.01	0.13	
	(0.10)	(0.17)	(0.15)	
Observations	1,789,042	1,789,042	1,789,042	1,789,042
Time Fixed Effect		✓	✓	✓
State Fixed Effect			✓	
Time-Gap Fixed Effect		✓	✓	✓

Notes: this table presents estimates regressing nominal job-level wages on unemployment. The dependent variable is quarterly percent growth in nominal posted wages, from the Burning Glass main sample. Wage growth is trimmed at the 1st and 99th percentiles. The regressors are the change in state-quarter unemployment, and the change interacted with an indicator for whether unemployment is decreasing, from the 2010-2020 LAUS, in percentage points. Standard errors are in parentheses, clustered by state. Column (1) presents estimates without controls. Column (2) adds time fixed effects, interacted with an indicator variable for whether unemployment is decreasing, as well as the interaction of time fixed effects and the gap between current and previous job postings fixed effects. Column (3) adds in state fixed effects. Column (4) presents estimates without asymmetries, by only including unemployment changes and time fixed effects as regressors, as in Column (1). Counts refer to the number of differenced observations.

Table 26: Controlling for lagged gap between postings

Panel A: Controlling for lag of gap

	(1)	(2)	(3)
ΔU_{st}	-0.041	-0.039	-0.037
	(0.016)	(0.016)	(0.015)
$\Delta U_{st} \times I(\Delta U_{st} < 0)$	-1.771	-1.763	-1.758
	(0.106)	(0.107)	(0.107)
Lagged Gap		0.049	
		(0.006)	
Observations	707,894	707,894	707,878
Time Fixed Effect	✓	✓	✓
Lagged gap FE			✓

Panel B: Correlation between gap and its lag

	Gap between postings
Previous gap between postings	0.093
	(0.004)
Observations	714,530
Time Fixed Effect	✓

Panel C: Job specificity of gap between postings

	Gap between postings
Observations	513,142
R-squared	0.41
Job Fixed Effect	✓

Notes: Panel A presents estimates regressing nominal job-level wage growth on unemployment changes. The dependent variable is the quarterly percent growth in nominal posted wages, from the Burning Glass main sample. Wage growth is trimmed at the 1st and 99th percentiles. The regressor is the change in state-quarter unemployment, from the 2010-2020 LAUS, in percentage points. Standard errors are in parentheses, clustered by state. Panel B shows the result for estimating a regression with the current gap between postings as the dependent variable and the lagged value of the gap between postings as the explanatory variable. Panel C displays the R-squared of the regression of the gap between postings on job fixed effects. In all panels, we limit our sample such that it only includes observations for which the lagged value of the gap between postings is not missing. Counts refer to the number of differenced observations.

C Data Appendix

C.1 Description of Burning Glass Algorithm

C.1.1 Classifying Job Titles

This subsection describes the algorithm that Burning Glass uses to extract job titles from the text of vacancy postings, Burning Glass' "CleanTitle" field. The wage posted in the vacancy is not used at any stage in this process. There is natural language processing model that segments the posting into different parts (such as company description, position description, and so on) and then finds the text containing the raw job title. The raw job title is then cleaned and standardized as follows. First, the title is cleaned by matching to an existing dataset of job titles maintained by Burning Glass. Second, there is a step to remove any location identifiers from the text of the job title (since these populate the separate location field). Third, there is a further step to standardize punctuation and other aspects of writing.

C.1.2 Measuring Length of Vacancy Posting

Burning Glass' algorithm checks daily for new vacancy postings, when a new vacancy is posted, it is assigned a start date. The end date is the minimum of (a) the first seen date + 60 days, or (b) the last seen date; we discard the vacancy length information for vacancies that record a length of 60 days. The last seen date is the date that the posting is either (i) removed from the internet; or (ii) includes some text saying it is no longer available, not accepting applications. Burning Glass' algorithm checks daily if (i) or (ii) is true.

C.2 Representativeness of the Main Sample

This section expands on the discussion in subsection 2.1. We show that our main Burning Glass Sample is broadly representative of the population of US employment and establishments; with notable caveats that we will make clear.

Appendix Figure 6 shows that the main regression sample is broadly representative at the regional level. We calculate the share of vacancies in each state over 2010-2020, for our main sample in Burning Glass; and from the official source, the Bureau of Labor Statistics' Job Opening and Labor Turnover Survey. Overall, the distribution of vacancies at the state level match across the two datasets.

Appendix Figure 7 shows that the main regression sample is also mostly representative at the broad occupation level. We calculate the share of employment over 2010-2020 for the main sample in Burning Glass; and from official data, namely, the Occupational Employment Statistics (OES) for 2010-2020. The distribution of jobs across occupations is broadly similar across the two sources. The main discrepancy is in transportation, which is heavily overweighted in Burning Glass. One innocuous reason for the difference could be that transporation is an occupation with high turnover; so, will be over-represented in vacancies relative to employment.

Appendix Table 20 shows that the main sample is also reasonably representative at the detailed occupation level. We construct the employment share within each 6 digit SOC occupation from the OES, and corresponding shares for vacancies from the main Burning Glass sample. We regress the Burning Glass share on the OES share. Column (1) reports a coefficient of 0.84—meaning, an occupation with a 1 pp higher share in the OES has a 0.84 pp higher share in Burning Glass, suggesting relatively similar representation. The coefficient is similar in Columns (2) and (3), which add 2 and 3 digit occupation fixed effects, respectively. Indeed, column (3) has a regression coefficient of 0.92, suggesting that detailed occupation shares in Burning Glass and the OES are very similar.

Appendix Table 21 shows that the main sample is also reasonably representative at the county level. We construct log employment at the county level, from the 2016 County Business Patterns (CBP), and construct log vacancies at the county level for the Burning Glass main sample. We regress Burning Glass log vacancies on CBP log employment. Column (1) reports a coefficient of 0.98—meaning, county with 1 percent higher employment in the CBP has 1 percent higher vacancies in Burning Glass, suggesting relatively similar representation. The coefficient is similar in Columns (2) and (3), which add census division and state fixed effects, respectively.

Appendix Table 22 shows that selection is not cyclical—neither selection into online vacancies in general, nor selection into the main regression sample. In Panel A, the outcome variable is the quarter by state change in the full set of vacancies in Burning Glass, relative to JOLTS. Here, the denominator is all online vacancies posted in the state and quarter; whereas the numerator is the number of vacancies in the state and quarter from official sources. The regressor is the change in quarter by state unemployment. Column (1) has a regression coefficient of 0.1, meaning that a 1 pp increase in unemployment leads to a statistically and economically insignificant decline in ratio of online vacancies to total vacancies in the state. Column (2) reports a similarly small estimate, after added in state fixed effects. Columns (3) and (4) show similarly small numbers when weighting by total vacancies in each state, instead of employment. Panel B instead asks whether there is cyclical selection within Burning Glass, between the full set of online vacancies and the main sample. Specifically, the outcome is the quarter by state change in the number of vacancies in the main sample, relative to the full set of Burning Glass vacancies posted in the state and quarter. The estimates in Panel

B are tiny, suggesting no cyclical selection. For example, column (1) reports an estimate of 0.02, meaning that a 1 pp increase in state unemployment leads to a 0.02 pp increase in vacancies in the main sample, relative to the full set.

Next, we show that our main sample is also broadly representative of the population of establishments. To use information on establishments, we merge the Burning Glass main sample to establishment level information from Dun & Bradstreet. Section 6.1 describes the Dun & Bradstreet dataset, and the merge to Burning Glass, in detail. Briefly, Dun & Bradstreet is a business analytics company that collect information on the universe of establishments in the United States from 1990-2020. D&B collects data on employment at the start of the year, industry classification and establishment age. For cost reasons, we purchased data from D&B only on the largest 30% of employers in Burning Glass. We achieve a high merge rate—75% of our main sample matches to establishments in Dun & Bradstreet. Dun & Bradstreet is known to measure employment poorly for very small establishments. Once these establishments are removed, Dun & Bradstreet appears to match official employment sources at the industry and regional level (Haltiwanger et al, 2013). Our extract from Dun & Bradstreet does not include very small establishments, suggesting our extract should measure establishment employment reasonably well.

Appendix Figure 8 shows that state level coverage of establishments is similar in the Burning Glass main sample and in official sources. Appendix Figure 9 shows the main sample is mostly representative of the distribution of establishments across industries—other than being somewhat under-weight healthcare and over-weight retail. Appendix Table 24, Panel A shows that the distribution of establishment age is similar in the main sample and in official sources. Panel B shows that our merged dataset is under-weight small establishments but otherwise representative—probably because the data we purchased from D&B excludes small establishments. Panel C shows that standard measures of the idiosyncratic shocks facing establishments—namely job creation and job destruction as in Davis, Haltiwanger, and Schuh (1998)—are similar in the two datasets, suggesting that our D&B subsample adequately measures establishment outcomes at annual frequency.

C.3 Hazard Estimation of the Frequency of Wage Change

This subsection describes the procedure for estimating the hazard rate of the latent wage change. 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 then

$$L = \prod_i \prod_t \left(1 - e^{-\lambda \gamma_{it}}\right)^{I_{it}} \left(e^{-\lambda \gamma_{it}}\right)^{1-I_{it}}.$$

The first order condition $\partial \log L / \partial \lambda$ implicitly defines the maximum likelihood estimate $\hat{\lambda}$ as

$$\sum_i \sum_t \frac{I_{it}\gamma_{it}}{e^{\hat{\lambda}\gamma_{it}} - 1} = \sum_i \sum_t \left(1 - I_{it}\right) \gamma_{it}.$$

With a hazard estimate in hand, we can calculate the other statistics as follows. 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. We can calculate the hazard rate of wage increase and decrease in an analogous way, and thereby calculate the probability of wage increase and wage decrease.

C.4 Construction of Census Region Level Price Measures

In this subsection, we describe how we construct measures of the consumer price index at the quarter by census division level. These measures were not available before this paper—the closest analogue is the state level inflation series of Hazell et al (2022), which is not available after 2017.

The Bureau of Labor Statistics (BLS) publishes quarterly census division level prices after 2018, but not before 2018. We now describe how we calculate census division level prices before 2018. The BLS reports MSA level inflation for the largest 20 metro areas in the United States. For census divisions in which MSA level inflation data is available from the BLS, we take the mean price level across MSAs within each census division to create a division level series for quarterly prices.

Then, we develop a procedure to calculate prices in census divisions for which MSA level inflation data is unavailable from the BLS. These census divisions contain only "mid-size" MSAs as defined by the BLS, instead of the "large" MSAs for which the BLS reports dedicated a dedicated inflation series. Therefore we assign to these census divisions the series reported by the BLS for inflation for mid-size cities from the corresponding census region. Recall that census divisions are a grouping of states into 9 groups, whereas census regions are a coarser grouping of states into 4 groups. As it turns out, we do not have to assign multiple census divisions to the same census region during this step.

We splice together our series for census division level inflation, with the BLS series, in the final quarter of 2017. We use the same procedure, separately, for the measures of census division prices including and excluding shelter prices.

C.5 Details on Dun & Bradstreet

C.5.1 Merge

We merge the main Burning Glass sample to Dun & Bradstreet as follows. First, we clean firm names in Burning Glass and Dun & Bradstreet in order to carry out the match. The firm name cleaning algorithm closely follows the algorithm developed in Hazell et al (2021), which uses a combination of standard cleaning procedures and a machine learning algorithm. We began with a list of (unclean) unique employer names in Dun & Bradstreet and Burning Glass. Then we clean both sets of employer names in the same way. We truncate employer names to 128 characters, and then we manually correct the names of some large employers, making use of code from the NBER Patent Data Project. We additionally stripped common words ("The", "Corp.", "Company", etc.), all non-alphanumeric punctuation, spacing, and capitalization. Next, we implemented a fuzzy matching algorithm, called *dedupe,* to create clusters of similar employer names. Dedupe makes use of a combination of squared edit distance comparisons subject to a confidence score threshold (which we chose to be 0.5, or 50% based on sample performance), as well as a small sample of names with manual labelling provided as training. For computational reasons, we employ blocking to limit the number of comparisons for each name to roughly 90 percent of each group of names sharing the first two letters. Within each cluster of names generated by dedupe, we set all names to that of the most common employer. Finally, we merge this crosswalk back on to the original (unclean) firm names and set the names to the new, cleaned versions to complete the process.

Then, we match Burning Glass establishments to Dun & Bradstreet establishments, based on City, County, State and cleaned firm name. Here, City is a field developed by Burning Glass, which roughly corresponds to a metropolitan statistical area—Dun & Bradstreet reports a similar field. In order to save money, we merge Burning Glass only to the largest 30% of establishments in Dun & Bradstreet. We are able to achieve a 75% match rate for the Burning Glass main sample.

C.5.2 Validating the D&B Dataset

To validate our measure of establishment employment, we compare to official data from the Business Dynamics Statistics (BDS), and find a reasonably close match. Appendix Figure 8 shows that state level coverage of establishments is similar in the Burning Glass main sample and the BDS. Appendix Figure 9 shows the main sample is mostly representative of the distribution of establishments across industries—other than being somewhat under-weight healthcare and over-weight retail. Appendix Table 24, Panel A shows that the distribution of establishment age is similar in the main sample and the BDS, though our extract is under-weight young establishments. Panel B shows that our merged dataset is under-weight small establishments but otherwise representative—probably because the data we purchased from D&B excludes small establishments. Panel C shows that standard measures of the idiosyncratic shocks facing establishments—namely job creation and job destruction as in Davis, Haltiwanger, and Schuh (1998)—are similar in the two datasets. This step is important because our selection correction will use establishment information to measure idiosyncratic shocks.

D Controlling for the Length of Time Between Vacancies

This section repeats the baseline results after controlling for the interaction of current time fixed effects and the length of time since the last posting. We also present scatter plots that condition on the length of time between vacancies. These controls alter the finding that wages are flexible upward—instead, wages do not respond much to falls in unemployment. The finding that wages are downwardly rigid is unaffected by these controls.

These results are important for the reader to digest. However the new covariates are likely a "bad control" and will lead to a form of selection bias if they are added to the baseline regression. Without the bad control, there could still be selection bias in the baseline regression. However as we discussed in Section 6, this selection bias is small.

The importance of these controls does perhaps suggest that trends in wage growth could confound our finding that wages are flexible upwards. Therefore we develop controls for trends in wages that are arguably unaffected by the bad control problem. With these controls, the baseline result is unaffected.

D.1 Results: Controlling for Length of Time Between Posts

We now show how the results change, after controlling for the interaction of current time fixed effects and the length of time since the last posting. We repeat the baseline regression equation (1). However, we add a control, $\gamma_{t,t-j}$, which interacts time fixed effects γ_t with an indicator for the number of quarters since the last posting. Appendix Table 25 reports the results. Now, the response of wage growth to increases in unemployment and to decreases in unemployment are both statistically and economically indistinguishable from zero. By contrast, in the baseline results, wages do not respond to increases in unemployment but respond significantly and strongly to decreases in unemployment. Therefore omitting the control $\gamma_{t,t-j}$ is important for the finding that wages are flexible upwards, although the result that wages are rigid downward is unaffected.

We next show the same result visually. We repeat the main scatter plot, Figure 1, visualizing the relationship between wage growth and unemployment changes. However, we condition on the length of time between vacancy postings. In Appendix Figure 10 we study the main scatter plot separately for vacancies with a gap of 1 through 6 quarters between postings, in panels (a)-(f) of the figure. From the figure, conditioning on the gap between postings, the relationship between wage growth and unemployment changes is linear and flat. Wages appear to respond neither to increases nor to decreases in unemployment.

D.2 Length of Time Between Posts as a Bad Control

These results are important because they show that a single control, the length of time between postings, alters the finding of upward wage flexibility. However the length of time between postings is a "bad control". This covariate will bias estimates of the relationship between wages and unemployment if it is included in the regression.

The time between postings is a bad control because it may be affected by the change in unemployment between postings. For instance, if there is a large rise in unemployment, then jobs might post vacancies less often. Therefore the time between postings is also an outcome that is caused by the regressor, which invalidates its use as a control. Angrist and Pischke (2009) provide a rubric to detect bad controls: *"[o]ne moral of the bad-control story is that when thinking about controls, timing matters. Variables measured before the variable of interest was determined are generally good controls. In particular, because these variables were determined before the variable of interest, they cannot themselves be outcomes in the causal nexus."* If we were to control for the gap between vacancies j, the converse would apply: j is determined at the same time as wage growth, meaning it is an outcome variable that cannot be used as a control.

Consider an example in which the bad control leads to bias. Suppose there are two types of jobs: high quality jobs, which face good idiosyncratic shocks and rising wages; and low quality jobs, which face bad idiosyncratic shocks and falling wages. Suppose that during contractions, only high quality jobs post with a 1 quarter gap between vacancies. During expansions, all jobs post with a 1 quarter gap between vacancies, including low quality jobs. Conditional on posting with a 1 quarter gap, there is an omitted variable—the composition of job quality is different during expansions versus contractions. This omitted variable biases the cyclicality of wages towards zero when studying only jobs posting with a 1 quarter gap. Wage growth is relatively high during contractions, since only high quality jobs post with a 1 quarter gap. Wage growth is relatively low during expansions, because there are both high and low quality jobs posting with a 1 quarter gap.

D.3 Selection Bias: Baseline Regression and Bad Controls

We now show formally that the bad control introduces a form of selection bias that is not present in the baseline regression. The baseline regression may also suffer from selection bias. Section 6 of the paper discusses the latter form of selection bias, and finds that it is small using a Heckman selection correction.

To start, let us review the selection bias of the baseline regression, and recap our finding that this selection bias was small. As we discussed in Section 6, selection bias may occur if jobs only post vacancies when the value to doing so is positive. If so, then then the regression relating wage growth to unemployment changes takes the form of equation (8), namely

$$E\left[\Delta_{t,t-j}\log w_{ist}|U_{st},U_{s,t-j},V_{ist}\geq 0,V_{is,t-j}\geq 0\right]=\alpha-\beta\Delta_{t,t-j}U_{st}+\gamma E\left[\Delta_{t,t-j}\phi_{ist}|V_{ist}\geq 0,V_{is,t-j}\geq 0\right]. \tag{9}$$

As we discussed in Section 6, the left hand side is the conditional expectation of wage growth between periods t and $t-j$ in which the job posts vacancies, $\Delta_{t,t-j}\log w_{ist}$. V_{ist} is the value of posting a job in state s at time t. Jobs only post vacancies if the value to doing so is positive, so that $V_{ist}\geq 0$. ϕ_{ist} is idiosyncratic shocks to wages. The conditional expectation of $\Delta_{t,t-j}\phi_{ist}$ is not in general zero because the value of posting depends on aggregating conditions, $\Delta_{t,t-j}U_{st}$. As a result there is selection bias. In Section 6, we correct for this selection bias using a standard Heckman estimator, and find that it is small. Therefore estimating equation (9) by ordinary least squares is approximately unbiased. As such, equation (9) is well approximated by the data generating process

$$\Delta_{t,t-j}\log w_{ist}=\alpha-\beta\Delta_{t,t-j}U_{st}+\gamma\Delta_{t,t-j}\phi_{ist}, \tag{10}$$

where $\Delta_{t,t-j}\phi_{ist}$ is conditionally mean-independent of $\Delta_{t,t-j}U_{st}$, so that the regression can be estimated by ordinary least squares with minimal bias.

Using the same regression framework, we now show that controlling for the length of time between posts introduces selection bias, even if it is not present in the baseline regression. Consider a selection equation that links the length of time j between posts to idiosyncratic and aggregate conditions. We have for all j' a series of equations

$$I\left(j_{ist}=j'\right)=\delta_{j'}+\mu_{j'}\Delta_{t,t-j'}U_{st}+\nu_{j'}\Delta_{t,t-j'}\phi_{ist}. \tag{11}$$

Here, $I\left(j_{ist}=j'\right)$ is an indicator variable equalling 1 if the length of time between posts is j', which depends on the change in unemployment $\Delta_{t,t-j'}U_{st}$, as well as the evolution of idiosyncratic shocks $\Delta_{t,t-j'}\phi_{ist}$. For instance, if μ_1 is negative, then jobs are less likely to post with a gap of one quarter when unemployment increases. Likewise, if μ_{10} is positive, then vacancies are more likely to post with a gap of ten quarters when unemployment increases.

Suppose we modify the baseline regression by controlling for the length of time between shocks. As such, we modify the baseline regression equation (10) to

$$\Delta_{t,t-j}\log w_{ist}=\alpha-\beta\Delta_{t,t-j}U_{st}+\sum_{j'=1}^{N}\pi_{j'}I\left(j_{ist}=j'\right)+\gamma\Delta_{t,t-j}\phi_{ist}, \tag{12}$$

where $I(j_{ist} = j')$ is a series of fixed effects for whether the gap between vacancies equals j', and j' ranges between 1 and N.

In the new equation with the "bad control", selection bias re-emerges and one cannot estimate β without bias by ordinary least squares. To see this point, take conditional expectations of equation (12), including conditioning on the event that $j = j'$. Then we arrive at

$$E\left[\Delta_{t,t-j} \log w_{ist} | \Delta_{t,t-j} U_{st}, j = j'\right] = \alpha - \beta \Delta_{t,t-j} U_{st} + \pi_j + \gamma E\left[\Delta_{t,t-j} \phi_{ist} | \Delta_{t,t-j} U_{st}, j = j'\right]. \tag{13}$$

The final term is correlated with unemployment changes $\Delta_{t,t-j} U_{st}$. Therefore there is omitted variable bias and one cannot estimate β by OLS. To see this point, taking conditional expectations and manipulating the selection equation (11) implies

$$E\left[\Delta_{t,t-j} \phi_{ist} | \Delta_{t,t-j} U_{st}, j_{ist} = j'\right] = \frac{1 - \delta_j - \mu_j \Delta_{t,t-j} U_{st}}{\nu_j}.$$

Therefore the final term in equation (13) is correlated with unemployment changes. This selection bias arises for the reason that we discussed in the previous subsection—due to how the "bad control" alters the expected value of idiosyncratic shocks. Even if the baseline regression equation (10) is little affected by selection bias, conditioning on the length of time introduces new selection bias.

D.4 Additional Concern: Trend Wage Growth

We have argued that controlling for the length of time between posts is a "bad control". However the importance of trend controls does perhaps suggest that trends in wage growth could confound our finding that wages are flexible upward. To understand the concern, consider two jobs, the first of which posted with a 1 quarter gap, and the second with a 4 quarter gap. Suppose that unemployment trends downward at a constant rate over this period, whereas wages trend upward. As a result, the second job experiences a 4 times larger fall in unemployment and a 4 times larger rise in wages than the first job. Comparing the two jobs suggests that wages increase as unemployment falls. In fact, trends in both variables explain the comovement. Adjusting for trends, wages might not be flexible upwards.

To rule out this concern, we cannot control for the length of time between vacancies, due to the bad control problem. What is required is a proxy for the length of time between vacancy postings that is not a bad control. We use as a proxy the lagged length of time between vacancy postings, where the lag is at the job level.

This proxy is appealing for two reasons. First, the lagged length of time between vacancies spans much of the same variation as the current length of time between vacancies. As such, controlling for the lagged length of time will absorb much of the variation associated with the length of time between vacancy posts. We show this point in Appendix Table 26, Panel B, by showing that the length of time between vacancy posts is significantly correlated with its lag. The reason is that the length of time between vacancy posts is to a large extent a fixed characteristic of the job posting the vacancy. Appendix Table 26, Panel C, confirms this point by regressing the length of time between posts on job fixed effects. The R squared of this regression is high: around 40%. Therefore fixed characteristics of a job explain a large share of the length of time between vacancy posts.

A second advantage of our proxy is that the length of time between vacancies in the past is predetermined with respect to current changes in unemployment. Being predetermined, the lagged length of time is less vulnerable to the bad control problem, as the previous quote by Angrist and Pischke (2009) suggests.

Controlling for the lagged length leads to a simple test for whether how much the length of time between posts matters, that is likely robust to the bad control issue. If adding the control does not affect β and δ, then estimates of wage cyclicality are not affected by the length of time between postings. Conversely, if adding the control changes estimates of β and δ, the length of time between vacancy postings is important and requires further investigation.

Appendix Table 26, Panel A reports the results from controlling for our proxy for the length of time between vacancies. We find that the wage cyclicality estimates, β and δ, change little, meaning the length of time between posts does not matter very much for wage cyclicality. In column (1) we report estimates of wage cyclicality for the baseline regression (1), on the sample for which lagged lengths of time between vacancies are available. In column (2) we add a linear control for the lagged gap. Estimates of wage cyclicality change little. In column (3) we instead saturate the regression, by including fixed effects for the lagged gap length interacted with time fixed effects. Again, the results change little. In this regression, we ignore other sources of selection bias. As we have discussed, Section 6 finds that other sources of selection bias are small.

E Job Composition and Variance of Wage Cyclicality Estimates

This section formally proves that job composition raises the variance of wage cyclicality estimates. As such, regressions that do not correct for job composition may lack the power to detect downward wage rigidity.

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

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

Proof. Summing regression equation (3) over i yields

$$\sum_i \nu_{ist} \Delta \log w_{ist} = \alpha + \gamma_t + \beta \Delta U_{st} + \delta_{\text{Job Level}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \varepsilon_{st} \tag{14}$$

where $\varepsilon_{st} = \sum_i \nu_{ist} \varepsilon_{ist}$. We can substitute equation (2) into equation (4) to rewrite the regression that uses average wages as

$$\sum_i \nu_{ist} \Delta \log w_{ist} + \sum_i \log w_{ist} \Delta \nu_{ist} = \bar{\alpha} + \bar{\gamma}_t + \bar{\beta} \Delta U_{st} + \delta_{\text{Average}} I\left[\Delta U_{st} < 0\right] \Delta U_{st} + \bar{\varepsilon}_{st}. \tag{15}$$

For notational simplicity, we can rewrite equation (14) as

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

and equation (15) 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 (14) and (15) 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

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

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

$$V\left[\hat{\bar{\mathbf{b}}}|\mathbf{x}_{st}\right] = V\left[\hat{\mathbf{b}}|\mathbf{x}_{st}\right] + \left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}'_{st}\right)^{-1}\frac{1}{(ST)^2}V\left[\sum_{s,t}\mathbf{x}_{st}u_{st}|\mathbf{x}_{st}\right]\left(\frac{1}{ST}\sum_{s,t}\mathbf{x}_{st}\mathbf{x}'_{st}\right)^{-1} \tag{16}$$

The second term in equation (16) 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 $V\left[\hat{\mathbf{b}}|\mathbf{x}_{st}\right]$.

F Selection Correction

This subsection explains how we implement the non-parametric Heckman estimator of Das et al (2003). We implement the selection correction in two steps as follows:

1. In the first step, we estimate the probability of vacancy posting p_{ist}, for a job i in a state s and quarter t. We estimate a regression of ξ_{ist} on $n_{is,t-1}$ and $n_{is,t-2}$, establishment employment at the start of the year and the previous year. Recall that ξ_{ist} is an indicator for whether the vacancy posts. We interact the regressors with state-by-time fixed effects in a third order polynomial series regression. Similarly, we estimate $p_{is,t-j}$ from a regression of $\xi_{is,t-j}$ on $n_{is,t-j-1}$ and $n_{is,t-j-2}$, interacted with state-by-time fixed effects in a third order series regression.

2. Then, we re-estimate our baseline regression equation. However, we include as extra regressors, a third order polynomial series regression in our estimates of p_{ist}, $p_{is,t-j}$ as well as $n_{is,t-1}$ and $n_{is,t-2}$.