

# Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data\*

John Grigsby

Erik Hurst

Ahu Yildirmaz

February 26, 2019

## Abstract

Using administrative payroll data from the largest U.S. payroll processing company, we document a series of new facts about nominal wage adjustments in the United States. The data allow us to define a worker's per-period base contract wage separately from other forms of compensation such as bonuses. We provide evidence that the extent to which base wages adjust is likely the appropriate concept of wage stickiness in many macro models. Nominal base wage declines are much rarer than previously thought with only 2% of job-stayers receiving a nominal base wage cut during a given year. However, accounting for shifts in nominal base wages of job-changers implies that *aggregate* nominal wages are more flexible than the nominal wages of job-stayers. In addition, we provide evidence that the flexibility of new hire base wages is similar to that of existing workers. Finally, nominal base wage adjustments are state-dependent: downward aggregate nominal wage adjustments were much more common during the Great Recession than in the subsequent recovery period. Throughout, we highlight differences in the adjustment patterns of base wages and of broader wage measures that include bonuses. Collectively, our results can be used to discipline models of nominal wage rigidity.

---

\*We thank Mark Aguiar, Susanto Basu, Steve Davis, Fatih Guvenen, John Haltiwanger, Jonathon Hazell, Pete Klenow, Alan Krueger, Marianna Kudlyak, Andre Kurmann, Alex Mas, Ben Schoefer, John Shea, Rob Shimer, Gary Solon and Joe Vavra as well as seminar participants at the 2018 American Economic Association meetings, Chicago, the European Central Bank, Maryland, Michigan, Minnesota, Princeton, the San Francisco Federal Reserve, USC, and Wharton for helpful comments. Per our use agreement, ADP approved the paper topic ex-ante. Additionally, ADP reviewed the paper prior to distribution with the sole focus of making sure that the paper did not release information that would compromise the privacy of their clients or would reveal proprietary information about the ADP business model. Authors' contact information: jgrigsby@uchicago.edu, erik.hurst@chicagobooth.edu. and Ahu.Yildirmaz@ADP.com.

**Online Appendix:**  
**“Aggregate Nominal Wage Adjustments: New  
Evidence from Administrative Payroll Data”  
(Not for Publication)**

## **Appendix A    Benchmarking ADP Data**

In this section, we benchmark the ADP data to various other data sources.

### **Appendix A.1    Demographics and Worker Tenure**

Table A1 shows some additional summary statistics for our ADP employee sample pooling across all years (column 1) and for selected individual years (columns 2-4). In particular, we show statistics for 2008 (our first year of data), 2012 (a middle year of data), and 2016 (our last year of data). The age, sex, and tenure distributions in our ADP sample matches well the age, sex, and tenure distributions of workers in nationally representative surveys such as the Current Population Survey (CPS). Additionally, according to the BLS, median tenure for workers over the age of 25 was about 65 months in 2012 and 2014 and was about 60 months in 2016. About one-fifth of our sample is paid weekly while three-quarters is paid bi-weekly. Less than five percent are paid monthly.

Given that ADP is growing over time, so too is our sample. Of our 1 million workers in our employee sample, only 220,000 are in our sample in 2008 while 377,000 are in our sample in 2016. Despite the growing sample size over time as ADP expands its business, the demographic composition of workers is essentially constant over time. One distinction is that average tenure is falling over time. Given that the Great Recession occurred early in our sample, it is not surprising that average tenure fell as many workers became displaced during the recession and eventually re-entered employment as the recovery took hold post-2012. Indeed, the roughly 6-month decline in worker tenure between 2012 and 2016 is also found in BLS data. However, worker tenure in the ADP data is higher in 2008 than similar 2008 numbers reported by the BLS.

### **Appendix A.2    Fraction Paid Hourly**

For our sample, 66 percent are paid hourly with the remaining 34 percent being classified as salaried workers. According to data from the CPS monthly supplements, only 57 percent of

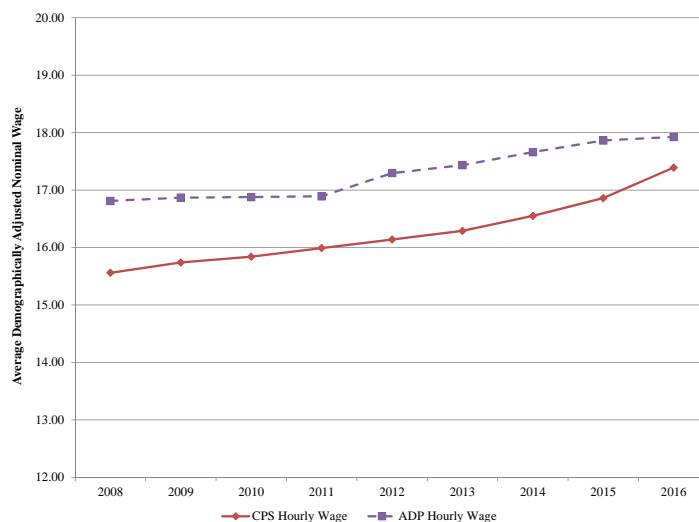
Table A1: Statistics for Employee Sample, Selected Years

	All	2008	2012	2016
Number of Workers	1,000,000	220,817	388,214	377,023
Number of Firms	91,577	37,269	52,478	45,519
Number of Observations	24,831,316	1,424,364	3,017,746	3,053,743
Age: 21-30 (%)	25.4	25.8	24.5	26.7
Age: 31-40 (%)	24.0	25.3	23.8	24.2
Age: 41-50 (%)	23.9	25.2	24.4	22.2
Age: 51-60 (%)	20.5	17.7	21.2	20.7
% Male	54.0	54.1	53.9	55.2
Average Tenure (months)	68.5	71.0	68.9	66.6
% Paid Weekly	20.7	21.4	21.7	21.0
% Paid Bi-Weekly/Semi-Monthly	76.0	75.5	75.1	75.1
% Paid Monthly	3.3	3.1	3.2	3.9
% Hourly	65.8	66.1	66.2	64.6

Note: Descriptive statistics for our employee sample in all years, 2008, 2012, and 2016. All data weighted to be representative of BDS firm size by industry distribution for firms with more than 50 employees. This table does *not* select on employees being between 21 and 60 years old

employed workers in the U.S. between the ages of 21 and 60 report being paid hourly during this time period. The difference between the CPS and ADP data may arise as the distinction between hourly and salaried workers is sometimes unclear within the ADP dataset. Some workers in the ADP data are automatically entered as having worked 40 hours each week at a given hourly wage. These workers are therefore classified as hourly wage workers. However, on many levels, these workers operate as if they were salaried: their actual hours never vary across weeks and not actually tracked in any meaningful sense. For these workers, their hourly contract wage is just their weekly salary divided by 40. Furthermore, these workers may report being salaried in survey data such as the CPS. For our purposes, however, we consider these workers as hourly, matching the ADP-provided definition. Additionally, with respect to wage changes, all changes in per-period earnings for these workers will be associated with a change in the hourly wage given that from the payroll system's perspective hours are fixed at 40 hours per week. Despite these differences in classification, the fraction reporting being paid hourly in the ADP data is broadly similar to the CPS averages.

Figure A1: Hourly Wage Comparison ADP vs. CPS, 2008-2016



Note: Figure shows the average hourly wage for hourly workers in our ADP sample and in a similarly defined sample of CPS respondents. Specifically, the CPS sample is restricted to workers between the ages of 21 and 60 who are paid hourly. For the average hourly wage for workers paid hourly in the CPS, we use data from the monthly outgoing rotation files from the CPS. In the outgoing rotation files, workers paid hourly are asked to report their hourly wage. ADP hourly wages reflect base wages. The ADP data is weighted so it is representative of the aggregate industry  $\times$  size distribution. The CPS data is weighted by the corresponding survey weights for the respective samples.

### Appendix A.3 CPS Comparison Average Hourly Wage for Hourly Workers

Figure A1 compares the average hourly base wages for hourly workers in our ADP sample to average hourly wages in a similarly defined sample of 21-60 year olds in the CPS. To get the hourly wage in the CPS, we use data from the outgoing rotation of respondents from the CPS monthly surveys. In the outgoing rotation, workers are asked if they are paid hourly and if so their hourly wage. For hourly workers, hourly wages are slightly higher in the ADP sample than in the CPS. This may be the result of the fact that, as discussed above, some salaried workers are classified as being hourly within the ADP data. Additionally, the ADP dataset does not include workers at small firms who are, on average, paid slightly less than workers at larger firms. The differences, however, between the ADP sample and the CPS sample are small and the trends are very similar suggesting that the ADP data is roughly representative of the entire U.S. population.

## Appendix A.4 Distribution of Changes in Annual Earnings

Although our paper is the first to use large-scale administrative data to measure wage adjustment in the United States, we are not the first to consider fluctuations in labor earnings. In particular, Guvenen et al. (2015) estimate a life cycle earnings process using earnings records provided by the Social Security Administration (SSA). Although their dataset has no measure of hours nor any breakdown of earnings by type (precluding a study of wage rigidity or the influence of bonus pay), it has the advantage of covering the universe of American workers over a long time span. As a result, their data is free of any form of sample selection, and represents a logical benchmark for our ADP dataset. We benchmark our annual earnings changes to Guvenen et al. (2015)’s Figure 1, which plots the distribution of individuals’ log annual earnings changes between 1998 and 1997, a period of relative calm in the labor market. Since our data do not extend back to 1997, we consider a year we deem relatively similar within our time period - the recovery years of 2015-2016.

The central challenge in benchmarking our data to annual earnings records is that the ADP data do not follow a worker if they move to employers who are not clients of ADP. This can generate large swings, both positive and negative, in annual earnings, which are not observed in datasets with the universe of employment, such as the SSA. Furthermore, since a great deal of annual earnings fluctuations arise from employment transitions, simply conditioning on workers appearing in the ADP data for a full 12 months will also lead to inaccurate fluctuations in annual earnings.

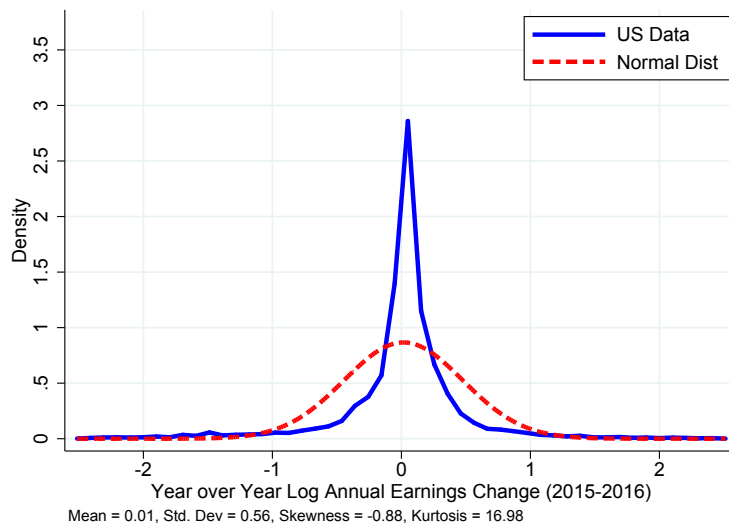
Our approach is somewhere in between the two extremes of treating all worker-years equally, and considering only full-year employment, in that we consider workers who appear in the ADP data for approximately the same number of months in both 2015 and 2016. Specifically, let  $N_i$  be the number of months that worker  $i$  appears in the ADP data in 2015. We consider only the annual earnings changes for workers who appear between  $N_i - x$  and  $N_i + x$  months in 2016, where  $x$  is a parameter that we set to 3 by default. For example, a worker who appears in the ADP data for 10 months in 2015 must appear in the ADP data for 7 to 12 months in 2016.

Figure A2 plots the distribution of log annual earnings changes in the ADP data. The figure matches the SSA data well but imperfectly. We estimate a mean earnings change of 0.01, in line with Guvenen et al. (2015). The standard deviation of annual earnings changes is 0.56 (compared with 0.51 in the SSA), while the skewness is -0.88 (vs -1.07) and kurtosis is 16.98 (vs 14.93).<sup>41</sup> We interpret these small differences to be the result of

---

<sup>41</sup>Varying  $x$  from 2 up to 4 does not have substantial impact on the results, but increasing  $x$  above 4 or decreasing it below 2 reduces the similarity between the ADP and the SSA data - the implied kurtosis is decreasing in  $x$  and standard deviation is increasing.

Figure A2: Comparison of Annual Earnings Changes in ADP Data with SSA Earnings data



*Notes:* This figure plots the distribution of year-over-year annual earnings changes for workers in the ADP data between 2015 and 2016. We limit attention to workers who appear in 2015 and 2016 for the same number of months, plus or minus 3. The blue line plots the realized distribution in the ADP dataset, while the red dashed line plots the normal distribution implied by the mean and variance of annual earnings changes. ADP data are weighted to reflect the aggregate firm size  $\times$  industry mix.

imperfectly capturing the annual earnings changes for job-changers, as well as transitions to unemployment. Overall, however, we find the similarity of this figure to that in Guvenen et al. (2015) to be encouraging.

## Appendix B Calculating Compensation Measures

This section details our construction of relevant compensation measures.

### Appendix B.1 Base Wages

The ADP data show an employee’s per period base payment rate. This administratively-recorded variable indicates the amount that an individual is contracted to earn every pay period. For hourly workers, this is literally an individual’s hourly wage, while it represents a salaried worker’s payment every week, if paid weekly, or every two weeks if paid biweekly, etc. Although these variables are administratively recorded, some employees still appear to have occasional errors in them, presumably resulting from keystroke errors. To deal with these issues, we clean the data in four ways. First, we code salaried workers who earn less than \$100 per pay period and have meaningful variation in hours worked as hourly workers. Second, we

winsorize wage rates below the federal minimum wage for service workers who receive tips. Some of these individuals may be unpaid interns who receive, for instance, transportation benefits from their employer. Third, we drop employees whose status codes indicate that their employment has been terminated. Finally, in our base wage change analysis we exclude workers who remain on the job but transition between being hourly and salaried workers. To compare the wages of hourly and salaried workers, we make the assumption that all salaried workers work 40 hours per week. This assumption does not affect our wage change calculations given that we exclude workers who transition from hourly to salaried status or vice versa; however, it is worth bearing in mind when we present statistics by employee wage percentile.

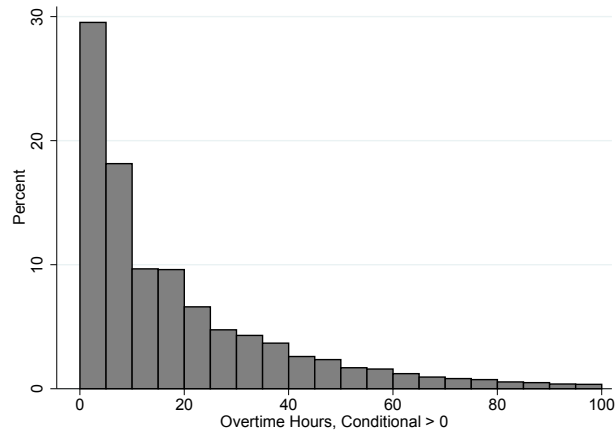
## Appendix B.2 Overtime Pay

In addition to the base payment rate and gross earnings variables, the ADP data include four separate earnings variables and four separate hours variables, denoting subcategories of compensation. These earnings and hours variables represent base earnings, overtime pay, or some combination of the two. These variables are not required for ADP clients to input. As a result, their quality and coverage is not comparable to that of gross earnings or base per period payment rates. Nevertheless, we use these variables to attempt to distinguish between overtime pay where possible.<sup>42</sup> To do so, we infer overtime premiums implied by these variables. For instance, we calculate implied base wages as base earnings divided by base hours, and overtime wage as overtime pay divided by overtime hours. The ratio of these implied wage rates to the administratively-recorded base wage provides a check on the validity of these implied wages. Most implied base wages, for instance, are exactly equal to the administrative base wage, and almost all lie between 1 and 1.1 times the contract wage. Overtime wage rates have large mass points at 1, 1.5, and 2 times contract wages: 75.7% of hourly workers with overtime premiums have implied wage rates which are 1.4-1.6 or 1.9-2.1 times their base wage. If the overtime wage is no more than 1.1 times the base contract wage, we declare overtime earnings to be part of base pay - although the worker may have worked overtime, she did not see any increased wage as a result, and so overtime cannot be a source of wage adjustment. Next, suppose the overtime wage is equal to  $1+x$  times the base wage, where  $x$  is at least 0.1. In this case, we define the overtime pay to be  $x$  times the number of overtime hours reported. This is therefore the *additional* pay, over and above what would be implied by their base wage.

---

<sup>42</sup>When available, the sum of these four earnings variables plus a variable defined as “earnings not related to hours” is always equal to the administratively-recorded gross earnings variables in our sample. The composition of earnings type across these five earnings categories is measured with error.

Figure A3: The Distribution of Monthly Overtime Hours



Note: Figure shows the distribution of monthly overtime hours, conditional on any overtime hours, for workers paid hourly in our employee sample.

After this imputation, 67.9% of worker-months record zero overtime. The distribution of monthly overtime hours, conditional on working any overtime, is presented in Figure A3. The figure shows that half of all overtime workers have less than 10 hours of overtime work per month. As a result, our focus on base pay and bonuses is likely sufficient to account for the majority of meaningful wage adjustments in the economy.

### Appendix B.3 Fringe Benefits

Fringe benefits were not required to be reported on a person's paycheck until 2012. As a result, our fringe benefit information is only reliable from 2012 onwards. We focus on the two largest forms of fringe benefits - health insurance and pension benefits. We define health insurance to be the sum of employer-provided health and accident plans, employer contributions to Health Savings Accounts (HSAs), Medical Spending Accounts (MSAs), and other Section 125 Medical Cafeteria plans. Pension benefits are the sum of employer contributions to 403(b), 501(c), 414(h), 401(k), 408(p), 408(k), and Roth 457 plans.



## Appendix C Robustness of Nominal Wage Adjustments for Job-Stayers

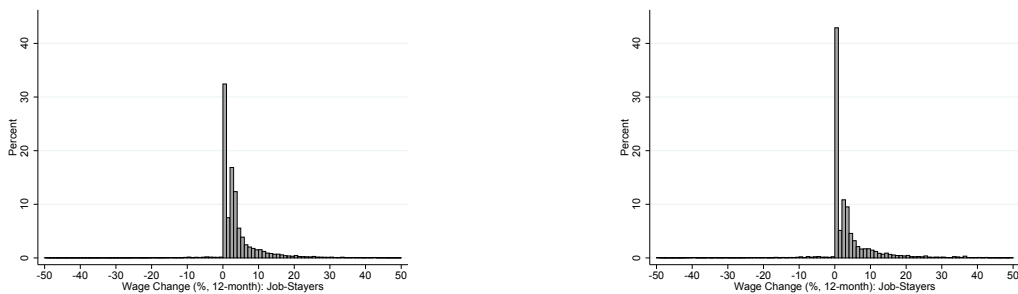
### Appendix C.1 Similarity in Patterns across Compensation Arrangements

The patterns of nominal wage adjustments for job-stayers are fairly robust across workers who are compensated in different ways. The top panel of Figure A4 shows the patterns of nominal base wage adjustments separately for non-commission workers (left) and commission workers (right). The bottom panel shows similar patterns for non-commission workers who do not receive a bonus (left) and non-commission workers who do receive a bonus (right). All of those panels pool together hourly and salaried workers. The patterns are strikingly similar across the four groups. Notice that essentially none of the groups receive a nominal cut to their base wage. All groups have between 30 and 40 percent of workers receiving no nominal base wage adjustments during the 12 month period. Non-commission workers who receive an annual bonus are the *most* likely to get a nominal base wage increase during the year. These workers both receive a bonus and are more likely to receive a wage increase. As seen above, these workers are more likely to be high earning workers. Conversely, roughly 40 percent of commission workers receive no nominal base wage change during the year. Finally, the patterns of nominal base wage adjustment for workers who receive essentially all of their earnings from base pay – non-commission workers who do not receive a bonus – are nearly identical to the patterns for all workers highlighted in Figure 2.

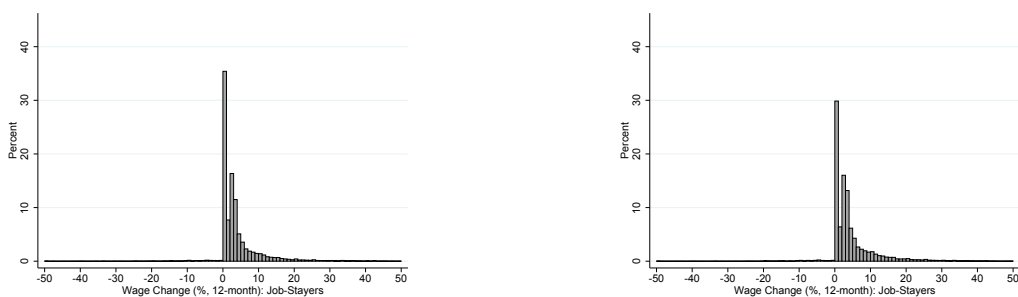
Figure A5 explores the extent to which nominal base wages are allocative. Specifically, we focus on our sample of hourly workers whose monthly hours worked fluctuates over the year. The number of pay weeks in the month varies over time, so we adjust our monthly hours for the number of pay periods making a measure of hours worked per week. We restrict the sample to only include those households whose hours worked per week varies substantively over the year.

The left hand panel of Figure A5 shows that wages are potentially allocative for these workers. Exploiting the panel nature of the data, we show that one-year base wage changes are associated with one-year hours worked changes, with an elasticity of 0.23. The right hand panel of the figure shows the one-year distribution of nominal base wage changes. It is nearly identical to the results shown in Figures 2 of the main text and A4. Even for workers whose hours fluctuate, there are essentially no nominal base wage cuts and roughly one-third of workers do not receive a year-over-year nominal base wage increase.

Figure A4: 12-month Changes in December Base Wages, 24-month Job-Stayers



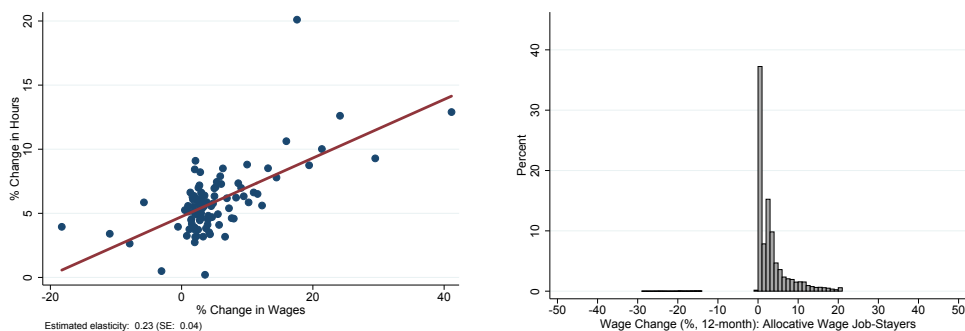
PANEL A: NON-COMMISSION IN YEAR  $t - 1$     PANEL B: WITH COMMISSION IN YEAR  $t - 1$



PANEL C: NO BONUS IN YEAR  $t - 1$     PANEL D: WITH BONUS IN YEAR  $t - 1$

*Notes:* Figure plots the 12-month change in December contract wages between year  $t - 1$  and  $t$  for workers who remain on a job for at least 24 months. Panel A plots the distribution of changes for workers who do not work commission in year  $t - 1$ , while Panel B plots the distribution for commission workers in  $t - 1$ . Panels C and D plot the distribution for workers who did and did not receive a bonus in year  $t - 1$ , respectively, excluding workers who work for commission.

Figure A5: 12-month Base Wage Changes, Job-Stayers, Hourly Workers w/ Variable Hours



PANEL A: CHANGES IN HOURS VS CHANGES IN WAGES    PANEL B: BASE WAGE CHANGE DISTRIBUTION

*Notes:* Figure shows results from a sub-sample of hourly workers whose weekly hours varies over the year and who remained continuously employed with the same firm during the 12 month period. We pool results over the entire 2008-2016 period. The left hand picture shows the relationship between the percent change in nominal base wages over the 12 months and the percent changes in hours worked. Each dot is a percentile of the wage change distribution. The right panel shows the distribution of the 12-month nominal base wage change.

Table A2: Base Wage Change Statistics, Pooled 2008-2016 Sample of Job-Stayers

	Monthly	Quarterly	Annual
<u>Unconditional</u>			
Skewness of Base Wage Changes (%)	9.7	5.3	2.8
Kurtosis of Base Wage Changes (%)	175.4	49.5	14.4
<u>Conditional on Any Wage Change</u>			
Skewness of Base Wage Changes (%)	2.1	2.1	2.4
Kurtosis of Base Wage Changes (%)	15.9	13.6	12.1

Note: Table shows higher order moments of the wage change distribution for different horizons for a sample of job-stayers in the ADP data between 2008 and 2016. For this table, we use our employee sample and pool together hourly and salaried workers. All data are weighted to be nationally representative of sample of workers working in firms with more than 50 employees.

## Appendix C.2 Higher Order Moments of the Base Wage Change Distribution

Table A2 shows higher order moments of the base wage change distribution for job stayers. In particular, we highlight both the skewness and kurtosis of the unconditional and conditional base wage change distribution.

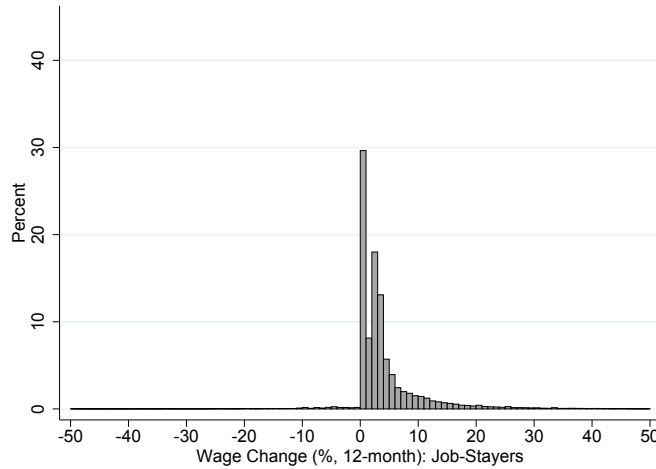
## Appendix C.3 Robustness to weighting and sampling

The analysis presented in the main text shows the distribution of wage changes for workers in large firms, weighted to match the firm size  $\times$  industry mix implied by the Census' BDS. Since a firm in the ADP data is defined by a unique ADP client, our weighting procedure may introduce bias if, for instance, large firms are especially likely to have multiple sub-units each of which separately contracts with ADP. To explore the potential bias, we show some of key results without any additional weighting.

Figure A6 plots the unweighted distribution of 12-month base wage changes for job-stayers. The only difference between this figure and Figure 2 of the main text is that here we do not weight data in order to match the firm size  $\times$  industry mix implied by the BDS. The patterns presented in this figure are almost identical to those in the main text, suggesting that our choice of weighting does not drive our results. While we only show this robustness for our base wage change results for job-stayers, the unweighted versions of other key results in the paper are also unchanged (e.g., bonuses, job-stayers, etc.).

The reason that our results are relatively insensitive to our weighting procedure is that the ADP data's firm size  $\times$  industry mix is fairly representative of the US economy, and

Figure A6: 12-month base wage change distributions for job-stayers: unweighted



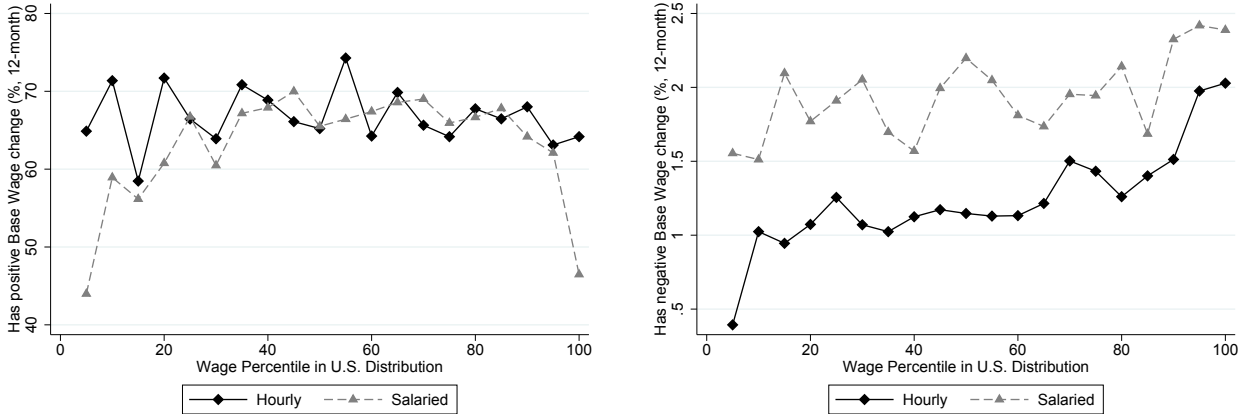
Note: Figure plots the unweighted distribution of 12-month base wage changes for job-stayers.

there are only relatively small differences in wage adjustment patterns across firm size and industry. We highlight this second fact in the next section.

## Appendix C.4 Heterogeneity in Base Wage Adjustment for Job-Stayers

Figure A7 plots the probability that a worker receives a year-over-year base wage change according to her initial position in the national base wage distribution. We calculate the wage distribution within hourly and salaried bins, and plot the patterns separately for each payment type. The black solid line shows the patterns for hourly workers, while the gray dashed line shows the patterns for salaried workers. The figure shows little systematic difference in the probability of receiving a base wage increase by initial wage percentile. However, those at the very top of the salaried distribution, for whom bonus income is a substantial portion of earnings, are less likely to receive an annual wage increase. Similarly, there is little relationship between wage level and the probability of receiving a base wage cut for salaried workers, with the probability of a wage cut bounded between 1.5% and 2.2% for much of the distribution. Low wage hourly workers, however, are less likely to receive a base wage cut than are high wage hourly workers, possibly due to minimum wage constraints or union contracts.

Figure A7: Probability of base wage adjustment by initial wage percentile, 2008-2016



PANEL A: POSITIVE CHANGES

PANEL B: NEGATIVE CHANGES

*Notes:* Figure shows the probability that a worker receives a year-over-year base wage increase (Panel A) or decrease (Panel B) by the worker’s initial position in the national wage distribution for workers of her payment type (i.e. hourly or salaried). This plot covers the period 2008-2016, and plots the patterns separately for hourly workers (black solid lines with diamond markers) and salaried workers (gray dashed lines with triangle markers).

## Appendix D Nominal Wage Adjustments for Job-Stayers by Firm Size and Industry

In this section, we document the extent to which wage adjustment varies by firm size. Additionally, we explore the potential bias in our key results from excluding firms with less than 50 employees from our analysis.

Figure A8 shows the distribution of annual wage changes over the 2008-2016 period by firm size and industry. The top panel shows patterns for hourly workers while the bottom patterns for salaried workers. The figure shows that base wage changes are monotonically increasing in firm size for both hourly and salaried workers. In a given 12-month period, 63.4% of hourly workers and 66.5% of salaried workers in firms with under 500 employees receive a base wage change. The comparable numbers for firms with 5000+ employees are 78.9% and 76.8%, respectively. These results complement the finding in the literature documenting that workers receive higher wages in larger firms (Brown and Medoff, 1989). Not only are workers in large firms receiving higher wages they also have a higher frequency of nominal base wage adjustments. All of the variation across firm size groups is in the propensity to receive a nominal base wage increase. While nominal base wage cuts are rare for all workers, there is no systematic variation in the propensity of a nominal base wage cut with firm size. While there are differences in nominal base wage adjustment across firm

size, the differences are relatively small. The small differences by firm size explains why our weighted results and unweighted results are so similar to each other.

Figure A8 also shows that there is some degree of heterogeneity across industries with respect to base wage changes. For example, both hourly and salaried workers in the manufacturing industry are much more likely to receive a base wage change than workers in construction during our sample period. This is in part due to the differential cyclical patterns of construction workers documented in section 10. Again, while there are some differences across industries in the extent of nominal base wage adjustments, the differences are quantitatively small so that our weighted and unweighted results are not that different from each other.

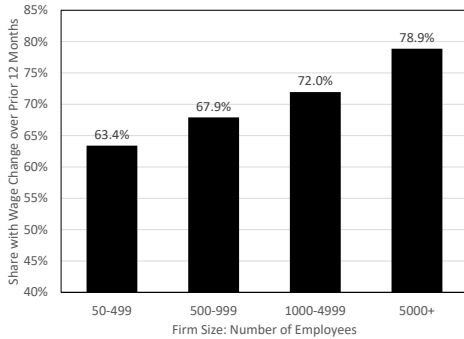
In order to further study the influence of excluding small firms with less than 50 employees from our baseline analysis, we use an additional dataset from ADP. This dataset originates from a payment product which is primarily marketed to firms with less than 50 employees. The dataset begins in June 2013 and contains similar measures of base wages and gross earnings to our main dataset that covers the 2008-2016 period for firms with more than 50 employees. Figure A9 plots the distribution of 12-month base wage changes for job-stayers in this small firm sample for the period 2014-2016.<sup>43</sup> The patterns for small firms are qualitatively similar to our patterns for mid size and larger firms - there remains a striking lack of wage cuts over a 12 month period among workers in small firms, as well as a substantial share of employees not receiving a wage change in a given year. Specifically, 48.7% of workers in small firms receive no wage change while 2.2% of workers receive a wage cut. As a reminder, the comparable numbers in firms with more than 50 employees were 34% and 2.4%. This findings are consistent with the results above that base wages adjust less frequently for workers in smaller firms.

How much can the exclusion of small firms from our main analysis bias our results? The BDS shows that 27.1% of workers were employed in small firms in 2016. We can use the results in Appendix Figure A9 to compute a new measure of the probability of nominal base wage adjustments inclusive of workers in small firms. Accounting for these small firms leads to a corrected annual probability of base wage change of 62.2% ( $0.271 \times 51.3\% + (1 - 0.271) \times 66.3\%$ ). Using the same procedure, the probability of a year-over-year base wage cut is 2.3% for all workers inclusive of those at small firms. Note, that these adjusted probabilities are

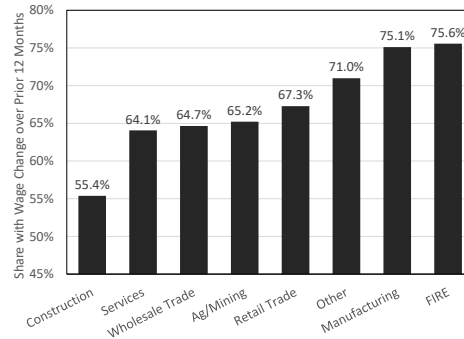
---

<sup>43</sup>We do not use this small firm sample in our main analysis for three reasons. First, this dataset does not contain any information on overtime, nor sufficient information needed to construct bonus payments reliably. Second, we are unable to track workers as they move from the small firm to the large firm sample. As a result, the rarity of job-changing from small ADP firm to another small ADP firm confounds our ability to measure wage adjustment of job-changers. Finally, the lack of a sufficiently long time series precludes the study of state dependence in wage adjustment in this dataset.

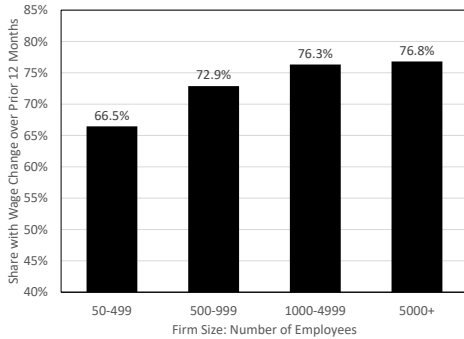
Figure A8: Share with Base Wage Change by Firm Size and Industry, All Years



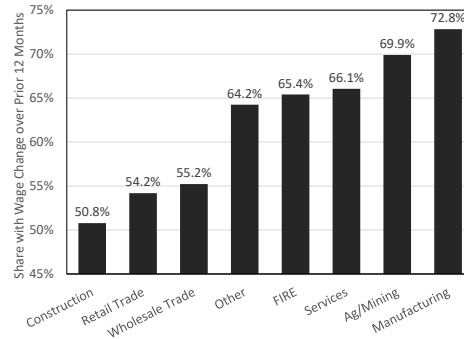
PANEL A: HOURLY WORKERS BY SIZE



PANEL B: HOURLY WORKERS BY INDUSTRY



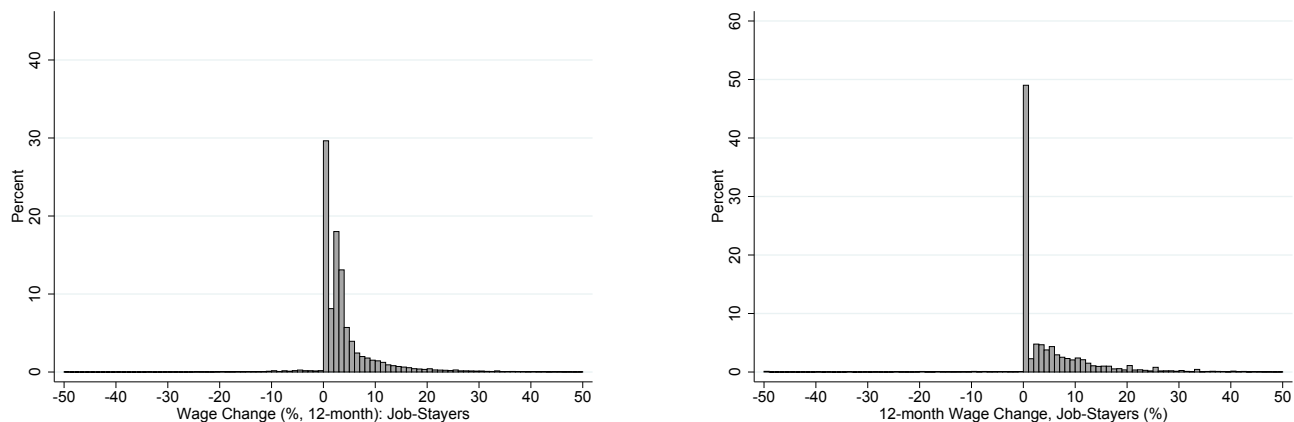
PANEL C: SALARIED WORKERS BY SIZE



PANEL D: SALARIED WORKERS BY INDUSTRY

Note: Figure shows the probability of receiving a base wage change by firm size and industry for our employee sample of job-stayers in the ADP data between 2008 and 2016. For this figure, we use our employee sample, and separately plot the patterns for hourly workers (Panels A and B) and salaried workers (Panels C and D). All data are weighted to be nationally representative of sample of workers working in firms with more than 50 employees.

Figure A9: 12-month base wage change distributions for job-stayers: firms with less than 50 employees



PANEL A:  $\geq 50$  EMPLOYEE FIRMS - UNWEIGHTED

PANEL B:  $< 50$  EMPLOYEE FIRMS

Figure plots the distribution of 12-month base wage changes for job-stayers for a sample of workers employed in firms with less than 50 employees. The results in this figure are produced using an ADP data product which covers the period Jan 2014 through December 2016 and only includes firms with less than 50 employees. The data in this figure are otherwise unweighted.

very close to those reported in the main text including only workers in firms with more than 50 employees (66.3% vs. 62.2% and 2.4% vs. 2.3%). To summarize, we conclude that the omission of workers in firms with less than 50 employees is not biasing our results substantively.

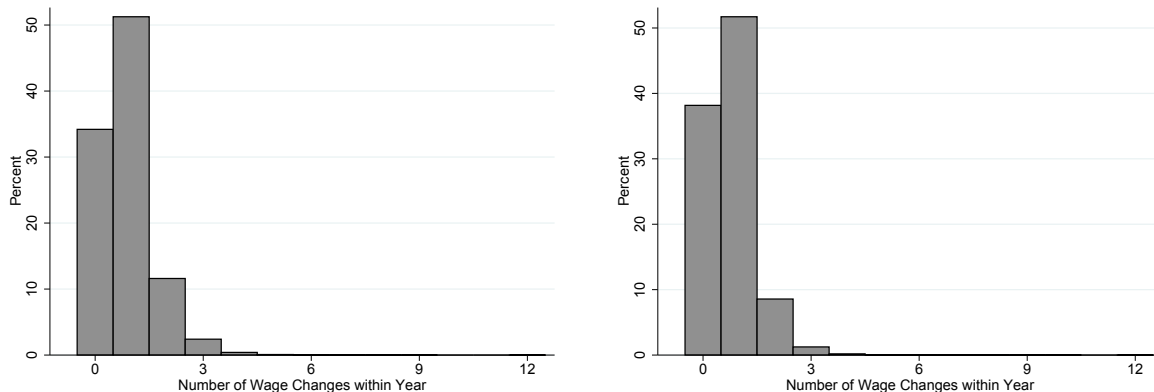
## Appendix E Time Dependence in Nominal Wage Adjustments, Job-Stayers

Many modern macro models assume some time dependence in wage setting. For example, Taylor (1979, 1980) emphasizes that staggered wage contracts can amplify business cycle persistence in response to aggregate shocks. New Keynesian macro models in the spirit of Christiano et al. (2005) use a Calvo (1983) model of wage setting. In this section of the appendix, we use our detailed micro data to explore evidence of time dependence in wage adjustment for our sample of job-stayers.

Figure A10 plots the average number of base wage changes during a given year for workers in our employee sample. As seen from Table 5 of the main text, roughly 35 percent of job-stayers receive no base wage change during a 12 month period. Over 50 percent of both hourly and salaried workers receive exactly one base wage change during a 12 month period when they remained continuously on the job. Between 10 and 15 percent of job-stayers



Figure A10: Number of Nominal Base Wage Changes over 12 month period, Job-Stayer Sample



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

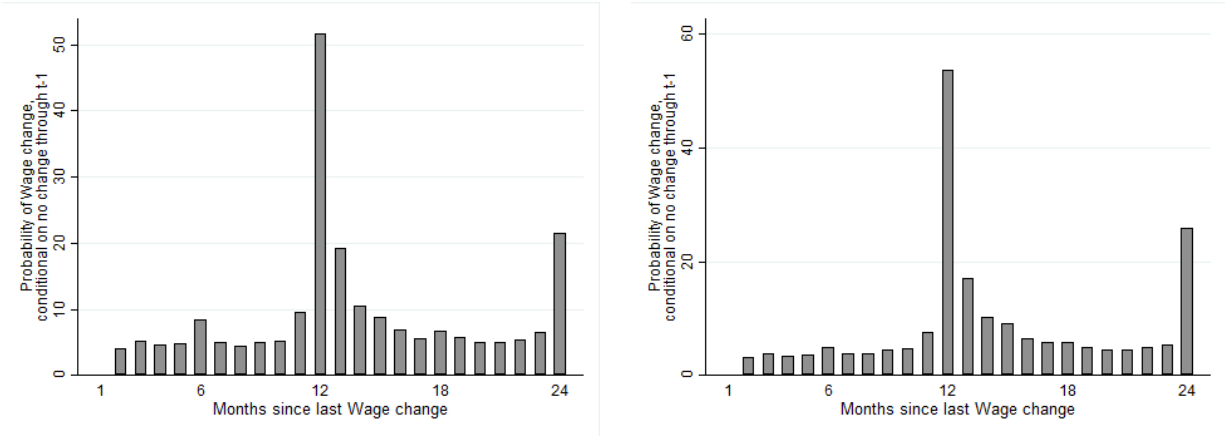
Note: Table shows the average number of nominal base wage changes for hourly workers (left panel) and salaried workers (right panel). We use our employee sample for this analysis and restrict our sample to those workers who remain continuously employed with the same firm during a 12 month calendar year. We use all data between 2008 and 2012 and average over the calendar years.

receive multiple base wage changes during a given year. The take away from Figure A10 is that roughly 90 percent of job-stayers receive either zero or one nominal base wage change during a given year. Multiple nominal base wage changes within a year are rare for continuing employees who remain on the same job.

To formally study time dependence in wage setting, Figure A11 plots the hazard functions of base wage adjustment for the subset of job-staying employees who experience at least two base wage changes over our sample period. Specifically, the figure shows the probability of a one month base wage change between  $t - 1$  and  $t$  conditional on the worker surviving to month  $t$  without a base wage change at the same firm.

The figure rejects the Calvo prediction that the probability of wage change is constant over time at the individual level for job-stayers. In most months, the probability of a base wage change is roughly constant at about 3-4%. However, roughly 12 months after the last wage increase, individuals are much more likely to get another base wage increase. Conditional on making it to month 11 with no base wage change, there is over a 50% probability that an individual gets a base wage increase in month 12. Note, given a little bit of calendar variation, there are small spikes at 11 and 13 months as well. We also see another spike in the hazard at 24 months. Moving away from a hazard analysis, we can define a sample of individuals who remained on their job for the next 18 months after a prior base wage change. We can then ask how many of these workers got their next base wage change 11-13 months later. Of consistently employed workers, 30% receive their next base wage change exactly

Figure A11: Hazard Function of Base Wage Change, Job-Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

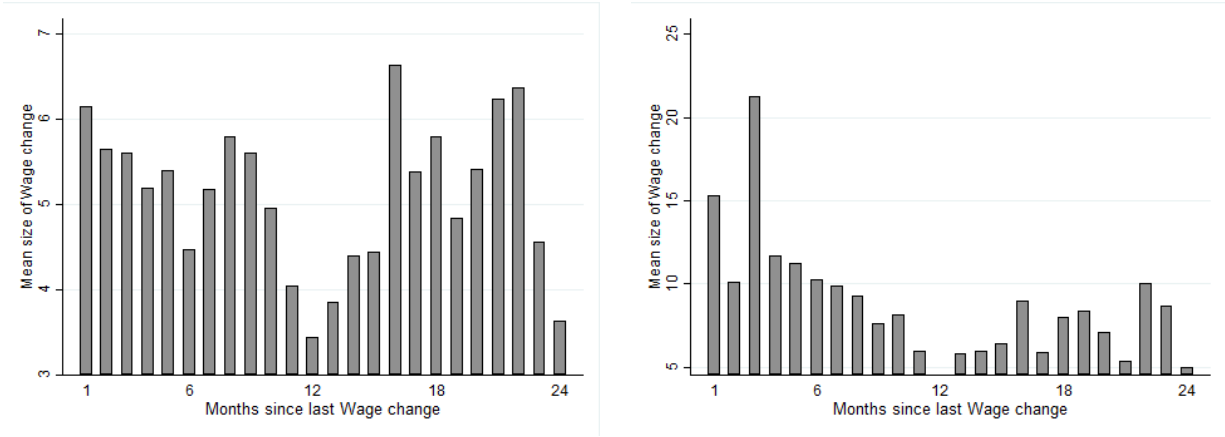
Note: Figure shows the hazard rate of a base wage change between  $t - 1$  and  $t$  conditional on surviving to  $t$  without a base wage change at the same firm. Sample only includes individuals with at least two base wage changes. We use all data between 2008 and 2016 for this analysis, and weight the data to be representative of the firm size  $\times$  industry mix in the BDS.

one year after their prior base wage change.

Figure A11 provides some evidence of time dependence in base wage adjustment. The majority of base wage changes occur annually. However, basic models of purely time dependent wage setting have predictions regarding the average size of wage changes. Under standard productivity processes with positive drift, individuals who are able to renegotiate their wage every month would negotiate smaller increases in their wages than those who renegotiate only once per year, on average. As a result, those who wait longer between wage changes should observe larger average changes in absolute value. We explore this prediction next.

Figure A12 shows the average size of the base wage change for job-stayers by the time since last base wage change. Since the vast majority of base wage changes for job-stayers are positive, this figure only includes workers who received a positive base wage change. While most base wage changes occur at 12 month frequencies, Figure A12 shows that the size of the base wage changes at these annual frequencies are much smaller than wage changes that occur at other times of the year. These predictions are not consistent with a standard Calvo (or Taylor) model at the individual level. However, the patterns could be consistent with a broader model of selection. If the workers who get these base wage changes that occur off-cycle are positively selected in some way, this could explain why they receive higher base wage increases. For example, if the worker receives an outside offer, the firm may have to raise the worker's base wage earlier than their annual cycle in order to retain the worker. Or,

Figure A12: Mean Size of Base Wage Changes by Time Since Last Change, Job-Stayers



PANEL A: HOURLY WORKERS

PANEL B: SALARIED WORKERS

Note: Figure shows the mean size of base wage increases for workers receiving a base wage increase  $t$  months after their last base wage change. Sample only includes individuals with at least two base wage changes. Additionally, we restrict our analysis to the job-stayer sample.

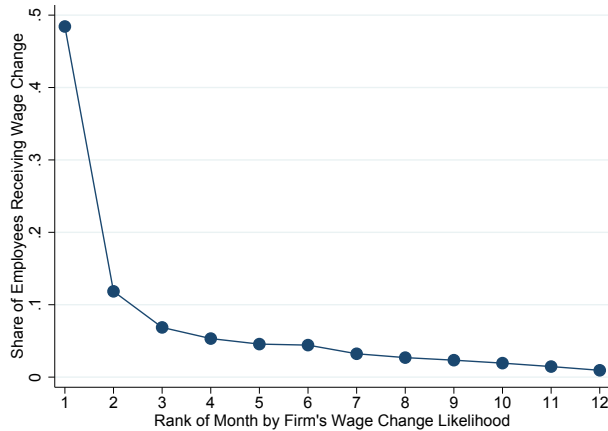
if a worker is promoted internally and the promotions are distributed throughout the year, it is not surprising that workers who receive a base wage change off cycle also get larger base wage changes.

Figure A13 shows the time dependence in wage setting at the firm level. For this analysis, we use our firm level sample. We restrict the firm-level sample to only include firms who remain in the sample of all 12 months during a given calendar year. Then, for each firm-year pair, we compute the fraction of workers who received a nominal base wage change during each calendar month. We then rank the months within a given firm-year pair from the month with the highest fraction of nominal base wage changes to the month with the lowest fraction of nominal base wage changes. For example, for some firms the highest month may be September while for other firms the highest month may be January. We then take the simple average probability of a worker receiving a base wage change across firm-year pairs for each ranked month.<sup>44</sup>

The figure shows that when a firm adjusts base wages, it tends to make all their base wage adjustments during one particular month of a given year. For example, a typical firm adjusts 50 percent of their workers base wages in the month where they make the most base wage changes. Given that only about 65 percent of workers get a base wage change (in the population as a whole) and the fact that we are averaging over firms and not workers,

<sup>44</sup>We also restrict our sample to only firm-year pairs where the firm adjusted at least 25 percent of their workers base wages at some point during the calendar year. This restriction is not too binding as 91% of firm-year pairs in our sample adjusted at least 25 percent of their workers base wages during the year.

Figure A13: Share Receiving Base Wage Change in Firm’s Months with Most Wage Changes, Firm-Level Data



Note: Figure uses data from our firm sample. We restrict the sample to include only firms who remain consistently in the sample during a given calendar year. We then compute for each calendar month within a firm-year pair, the fraction of workers who received a nominal base wage change during that month. We then rank the months within a given firm-year pair from highest month of base wage changes to lowest month of base wage changes. We then take the simple average across all firm-year pairs for each month rank. When making the figure, we restrict our analysis to only those firms who adjusted at least 25 percent of their workers base wages at some point during the calendar year.

the figure suggest that firms do most of their base wage changes in one month out of the year.<sup>45</sup> As a point of contrast, firms only adjust roughly 10 percent of their workers base wages in the second highest ranked month. The fact that the share of base wages adjusted are roughly flat between the second highest ranked month and the lowest ranked month is consistent with the worker data where some adjustments are occurring off-cycle at a roughly constant hazard. These changes are likely due to promotions and/or the response to outside offers.

While the Calvo predictions may be rejected at the individual and firm level, Calvo may still be a good approximation for the aggregate macro economy if firms stagger the months in which they adjust wages. Indeed, this is the underlying intuition behind the staggered wage contract model. Instead of each individual probabilistically getting a wage change each period, individuals deterministically get a wage change at a fixed frequency but a constant fraction of the wage contracts adjust each period. To see whether Calvo is a good approximation for job-stayers in the aggregate economy, we explore the extent to which base wage changes are coordinated within a given calendar month.

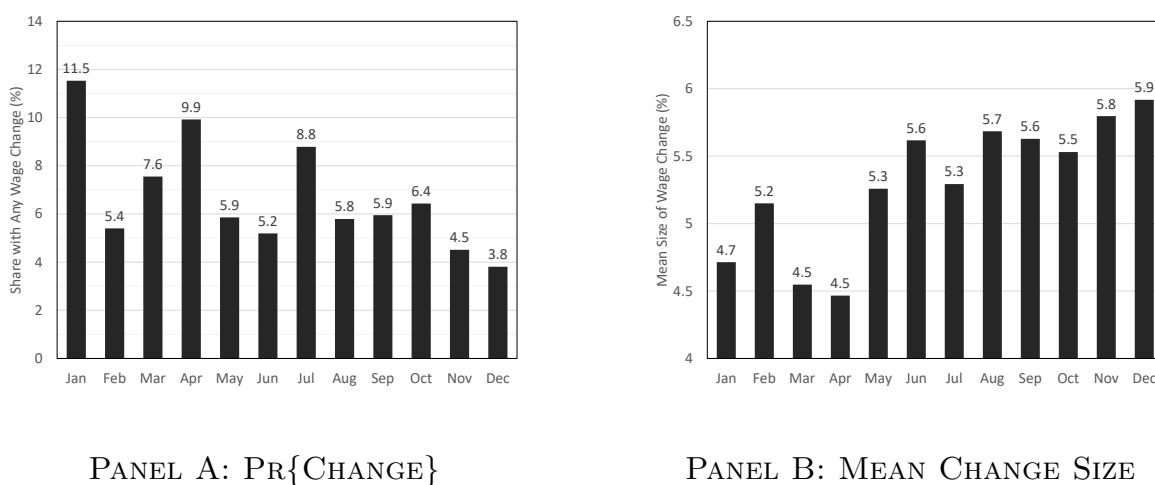
Figure A14 shows the probability of base wage changes by calendar month pooling to-

---

<sup>45</sup>This observation represents the labor market analogy to the price-setting rule employed in Midrigan (2011) in which multi-product firms enjoy economies of scale in coordinated output price adjustment.

gether hourly and salaried workers. For this analysis, we return to our employee sample and focus only job-stayers. The figure shows some slight seasonality in the data. The probability that a worker receives a base wage change is highest in January. The next highest months are the beginning months of each calendar quarter (April, July and October). However, these differences mostly wash out at the quarterly frequency: 23.4 percent of workers receive a base wage change in the first quarter of the year while 21.1 and 21.5 percent of workers receive a base wage change in the second and third quarters. Only 16.6 percent of workers receive a base wage change in the last quarter of the year.

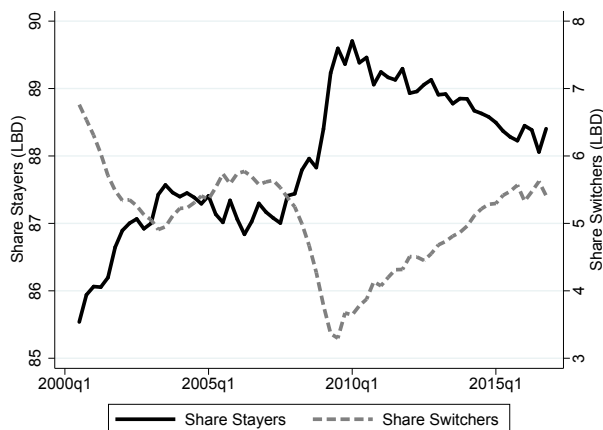
Figure A14: Seasonality in Base Wage Changes, Job-Stayers, All Years



Note: Figure plots moments of the base wage change distribution in each calendar month, averaged over our full sample of job-stayers pooled between 2008 and 2016. Panel A plots the probability of adjustment, while Panel B plots the mean size of a wage change, conditional on the change occurring. This figure combines hourly and salaried workers.

Overall, the evidence presented in this section shows strong evidence of time dependence in base wage adjustment. The majority of base wage changes occur annually, usually at the beginning of a firm’s fiscal year - either in January, April, or July. However, there is a roughly constant probability of base wage adjustment across the four quarters of the year, suggesting that models of Calvo adjustment may be a reasonable approximation of the base wage adjustment process. However, as we documented in Section 10 of the main text, base wage also adjustment appears to be state dependent. Modelers seeking to use a Calvo wage adjustment process should consider simple extensions, such as incorporating an asymmetric probability of base wage cuts and increases (see, e.g. Schmitt-Grohé and Uribe (2012)). Additionally, as we highlight throughout the text, annual bonuses also provide an additional marginal of flexibility for worker wages.

Figure A15: Time Series of the Share of Workers who are Stayers versus Switchers, LEHD Job-to-Job Flows data



Note: Figure plots the quarterly share of workers who are job-stayers (left axis, solid black line) and job switchers (right axis, dashed gray line) in the LEHD’s Job-to-Job (J2J) flows data over the period 2000Q1 through 2016Q4.

## Appendix F Time Series Trends in Aggregate Share of Job-Changers

Aggregate nominal base wage flexibility is a function of both the base wage adjustments for job-stayers and job-changers. Thus, measuring the cyclical nature of aggregate wage adjustments requires the evolution of the composition of job-stayers relative to job-switchers over the business cycle. Figure A15 shows the quarterly share of job-stayers and job-switchers between 2000 and 2015 from the Census’s Job-to-Job Flow Data (J2J), which is made from the LEHD. The difference between the sum of the two lines and one is the fraction of workers who left employment for longer non-employment spells during the quarter. During the Great Recession, the quarterly job-switching rate fell to 4 percent while during the 2012-2016 period the quarterly job-switching rate returned to a pre-recession level of about 5.1 percent. Job-staying rates were roughly the mirror image of job-changing rates. As above, we construct annual job changing rates by multiplying the quarterly rates by 4. Doing so implies that during the Great Recession 16 percent of workers switched jobs compared to roughly 20 percent during the recovery. We weight to match these proportions in the time series. Since job-changers receive nominal wage changes and cuts at a substantially higher rate than job-stayers, this composition effect pushes towards lower aggregate flexibility during the recession, even if both changers and stayers observe less downward rigidity in recession periods.