Online Appendix: "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data" by John Grigsby, Erik Hurst and Ahu Yildirmaz

Appendix A Data Details and Sample Construction

In the main text, we outlined a broad overview of our sample construction. In this appendix, we provide additional details.

Our principal sample is of one million random employees in the ADP database. To draw this sample, we first construct a list of all the unique IDs of employees who are between the age of 21 and 60 for at least one month in the data. Then we randomly select one million such IDs, and draw the complete history of each of these individuals in the data. By drawing their complete history, we track workers as they switch across employers, so long as they remain employed with a firm that contracts with ADP.

We construct a similar sample of 3,000 random firms. For this, we construct a list of all unique client codes in the data, and randomly select 3,000 such firms. We then draw all of the information for workers who work for those firms in each month. The first three rows of Table A1 show the number of unique employees, firms, and monthly observations in the employee and firm samples.

Approximately 0.4 percent of worker-firm observations appear multiple times in the same month. This principally results from workers who work multiple jobs within the same firm in that month, such as a line worker who is also a shift manager. The rarity of such events mean that they will not bias our estimates substantially. Since our focus is on the extent to which wages adjust, we only include the observation with the smallest base wage in the month. The intent is to look at fluctuations in wages for a worker in a given job.

For some analysis, we focus on individuals who switch firms. In order to maximize sample size for such analyses, we additionally draw the complete set of job-switches in the data. To construct this sample, we proceed iteratively. First, we bring in the complete data for May of 2008: the first month of our sample. Then we merge in the worker's wage and firm in May 2009 (if the worker only works for one employer in May 2009), and calculate the worker's wage change between the two years. We then do the same for June of 2008, and append this to our file. We continue to iteratively add months, keeping only the first and last month of a worker's employment with a particular firm.

Once we have a complete set of worker-firm employment spells and their associated 12-month wage changes, we first drop all firms who appear to undergo a merging of client codes as described below. We then drop one-month employment spells with a given firm, and workers who have overlapping employment spells at more than one firm; that is we restrict attention to single job-holders. This additionally removes about 8% of workers. We then keep only the last month of employment at each firm to arrive at our sample of job-changers.

When creating our job-changer sample, three additional issues are worth noting. First, we stress that we are measuring wage changes for workers who move from one ADP firm to another ADP firm. An implicit assumption we make throughout the paper is that the patterns of nominal base wage adjustments for workers who migrate across ADP firms are similar to the patterns of nominal base wage adjustment for workers who migrate to and from non-ADP firms.

Second, as noted above, multiple establishments within a firm sometimes contract separately with ADP or, on occasion, firms will spin off into multiple units each contracting separately with ADP. In this case, a movement from one establishment within a firm to another establishment within the same firm will look like a job-change. To account for such flows, we measure the percent of job-changers leaving a given firm in month t and showing up at another ADP firm in month $t+1$ or month $t+2$ using the universe of our data. If more than twenty percent of job-changers leaving firm i subsequently show up in firm j with no intervening employment spell elsewhere between t and $t+2$, we treat switches from i to j as within firm movements over this time period, and do not include them in our job-changer sample. In addition, if a worker's reported tenure does not reset after switching firms, we exclude that worker from the job-changer sample. Removing such dubious switches excludes around 20 percent of observed job changes. We also restrict our analysis to include only those workers who switch between either hourly jobs or who switch between salaried jobs. We exclude those who switch between the two types of jobs. These switches across payment types are relatively rare – only 5.9% of job switches involve such a change – but generate large swings in base wages in almost all cases. All of these restrictions leave us with 2.9 million unique job-switches with non-missing year-over-year wage changes.

Finally, the choice of timing with respect to job changes is more nuanced given the nature of our data. When we see a worker at firm i in month t and then see a worker at firm j in month $t+12$, the worker may have multiple other jobs in the interim. Because we only measure labor market outcomes for ADP firms, if a worker disappears from our dataset for a short time but reappears later, we are not able to distinguish if the worker was not employed or whether the worker was employed but at a non-ADP firm. For many applications, such distinctions are not important. However, it is worth keeping such timing issues in mind when interpreting our wage adjustment measures for job-changers. In some specifications we explore the sensitivity of our results to this timing issue by restricting our job-changing sample to only include individuals who left firm i in month t and appeared in firm j in month $t+1$. There are 1.1 million such switches in the data, approximately 37% of all observed job switches. For these workers, any intervening unemployment spell would be short; this can therefore be thought of as a sample of employed-to-employed (E-E) flows. This restriction does not change the conclusions of section **??**: we are unable to detect statistically significant differences in the cyclicality of wage changes for job-changers who have short gaps in ADP employment relative to those who have long gaps in employment. This is likely due to measurement error in our construction of E-E flows, as the existing literature has found that E-E flows drive much of the cyclicality in job-changers' wages (Gertler et al., 2016).

All replication codes from the paper are available on the AER website and the authors' websites. Given the confidential nature of the ADP data, we were only able to access the ADP data through APD provided laptops that directly accessed ADP servers. Given the sensitive nature of the data, we are not able to post any of the data online for replication purposes. However, replication is possible for researchers who obtain access to the ADP data. In terms of replication, we only use ADP's anonymous employee level database (as described in the text and in more detail above). The variable names in our code are the ones provided by ADP in their employee level data files. Researchers with access to the data can use our code directly with the ADP data to create all measures within our paper. Researchers interested in accessing the confidential ADP data can send requests to the ADP Research Institute (Ahu Yildirmaz at ahu.yildirmaz@adp.com). ADP Research Institute collaborates with researchers outside the institute, and through these research collaborations, researchers have the opportunity to conduct research based on ADP's anonymized, aggregated and de-identified payroll data.

Appendix B Benchmarking ADP Data

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

Appendix B.1 Firm Size Distribution

There are two areas of concern regarding the representativeness of the ADP data by firm size. First, the patterns we highlight in the paper apply only to firms with more than 50 employees. To the extent that the nature of nominal wage adjustments differs by firm size, the patterns we document within our sample may not be representative of the US economy as a whole. Below, in Appendix E, we show that there are only modest differences in nominal wage adjustments by firm size within our sample, we conjecture that any potential bias in our headline results from excluding firms with fewer than 50 employees is likely to be small. Furthermore, from 2013 onward, we also have access to ADP's data for firms with fewer than 50 employees. As we also highlight in Appendix E, these data reinforce that any potential bias from excluding small firms from the main results in our paper is likely to be minor.

The second concern is whether ADP clients are representative of firms with more than 50 employees. According to industry reports, roughly 50 percent of US firms in recent years report outsourcing their payroll services to payroll processing companies.[1] As noted in the main text, ADP processes payroll for about 20 million US workers per month. While ADP is the largest payroll processing company, the industry has many competing firms including Intuit, Workday, and Paychex.

According to these same industry surveys, very large firms (firms with more than 10,000 employees) are less likely to outsource their payroll functions. Appendix Table A1 highlights the employment-weighted firm size distribution in our "employee sample" (column 1) and in our "firm sample" (column 2). For the results in this table, we pool our data over the entire 2008-2016 period. By design, we randomly drew 1 million employees for our employee sample and 3,000 firms for our firm sample. Our employee sample includes roughly 250 thousand distinct firms while our firm sample includes roughly 3.3 million distinct employees. The number of actual observations is much larger for each sample because we observe employees for multiple months. For our employee sample, we track employees across all months between 2008 and 2016 that they are employed at *any* ADP firm. For our firm sample, we track all employees in that firm across all months that they remain employed at that firm.

For comparison, column 3 of Appendix Table A1 includes the firm size distribution from the U.S. Census's Business Dynamics Statistics (BDS) over the same time period restricting our attention to only firms with more than 50 employees.[2] As seen from the table and consistent with industry surveys, ADP under-represents very large employers (those with at least 5,000 employees). According to BDS data, nearly 46 percent of all employment in firms with more than 50 employees is in firms with more than 5,000 employees. The ADP data only have 19 percent of employment (in our employee sample) in firms with more than 5,000 employees.[3] As noted above, some of this difference also results from the fact that the ADP

[1] See, for example, https://www.feicanada.org/ajaxfilemanager/uploaded/RH_FERF_Benchmarking2014.pdf.

[2] According to BDS data, 72% of all U.S. employment during this time period is in firms with more than 50 employees.

[3] We also explore how the industry distribution of the ADP sample compares to the industry distribution in the BDS. We are unable to report ADP's precise industry distribution for disclosure reasons. The ADP sample has a slight over-representation amongst the manufacturing and broad service sectors, and a

Table A1: Firm Size Distribution in ADP Samples and the BDS, Pooled 2008-2016 Data

	ADP Employee Sample	ADP Firm Sample	BDS Data
Number of Employees	1,000,000	3,296,701	.
Number of Firms	254,729	3,000	.
Number of Observations	21,750,672	68,267,166	.
% Firm Size: 50-499	38.4	31.3	29.5
% Firm Size: 500-999	13.6	13.9	7.3
% Firm Size: 1000-4999	25.2	22.2	17.5
% Firm Size: $\geq$ 5000	19.0	32.5	45.6

Notes: Table reports the share of employees in firms of various sizes in our random samples of the ADP data, stratified at the employee (Column 1) and firm levels (column 2). Column 3 reports the associated employee-weighted firm size distribution reported in the Census' Business Dynamics Statistics (BDS) data. All numbers span the period 2008-2016. In addition, the first three rows show the number of unique employees, firms, and observations in each of our ADP subsamples.

definition of a firm is different from Census definitions.

To account for the concern that the data do not perfectly represent the universe of all U.S. firms with at least 50 employees, all analyses in the main text are weighted to match the BDS's firm size by industry mix of employment shares for firms with greater than 50 employees. We compute our weights for each year between 2008 and 2016. By re-weighting the data, we control for sample selection along these key observable dimensions. Although there may yet remain selection into the sample along unobservable dimensions, we consider these potential selection issues to be small once controlling for firm size and industrial mix.

Appendix B.2 Demographics and Worker Tenure

Appendix Table A2 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 match 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 211,000 are in our sample in 2008 while 366,000 are in our

complementary underweight in retail trade, construction, and agriculture.

Table A2: Statistics for Employee Sample, Selected Years

	All	2008	2012	2016
Number of Workers	1,000,000	210,573	373,168	365,739
Number of Firms	254,729	67,165	117,947	94,582
Number of Observations	23,382,167	1,335,656	2,841,901	2,891,002
Age: 21-30 (%)	25.2	25.2	24.4	26.6
Age: 31-40 (%)	24.0	25.4	23.8	24.2
Age: 41-50 (%)	24.1	25.6	24.4	22.3
Age: 51-60 (%)	20.6	18.0	21.2	20.8
% Male	53.9	54.1	53.7	55.0
Average Tenure (months)	69.2	71.8	69.5	67.2
% Paid Weekly	20.3	21.3	21.0	20.6
% Paid Bi-Weekly/Semi-Monthly	76.6	75.8	75.8	75.5
% Paid Monthly	3.2	3.0	3.1	3.9
% Hourly	65.4	65.4	66.0	64.5

Notes: 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

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 exception 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.

For our sample, 65 percent are paid hourly with the remaining 35 percent being classified as salaried workers. According to data from the CPS monthly supplements, only 57 percent of 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 reported hours never vary across weeks. 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.

Appendix B.3 CPS Comparison Average Hourly Wage for Hourly Workers

Appendix 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 B.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 how adjustment patterns differ by compensation component), it has the advantage of covering the universe of American workers over a long time span. As a result, their data are free of sample selection, and represent 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.

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

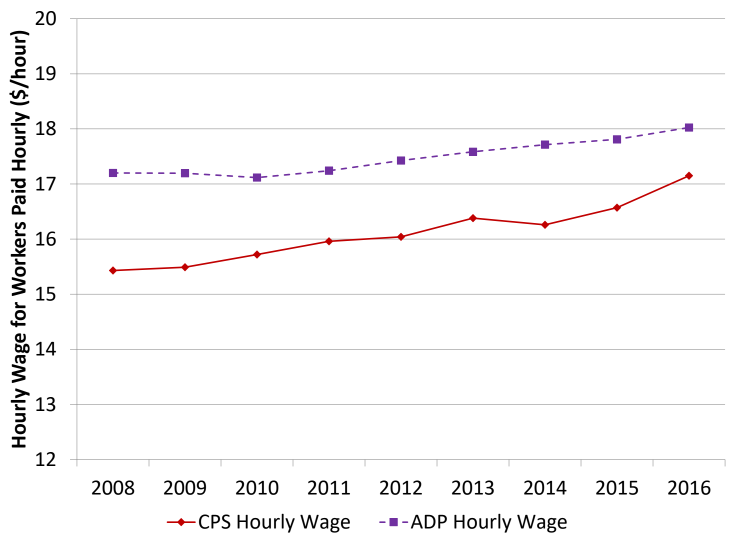

Notes: Figure shows the average hourly wage for hourly workers in our ADP sample and in a similarly defined sample of CPS respondents. Specifically, the CPS sample is restricted to workers between the ages of 21 and 60 who are paid hourly. For the average hourly wage for workers paid hourly in the CPS, 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 CPS data are weighted by the corresponding survey weights for the respective samples.

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.

Appendix 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

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

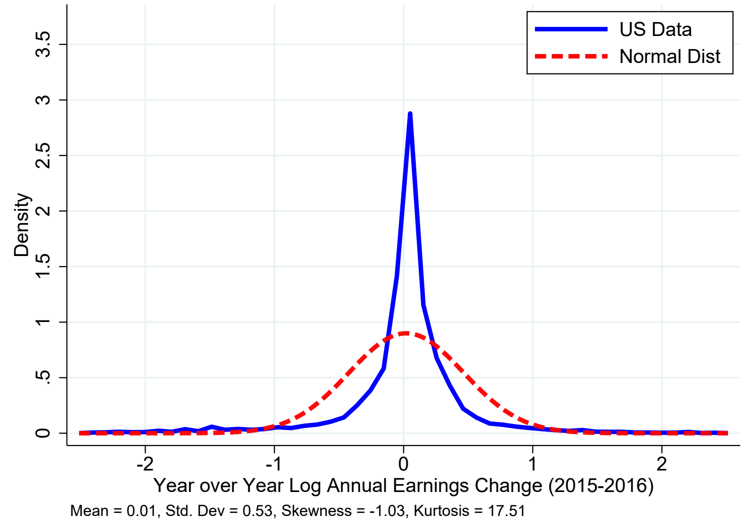

Notes: This figure plots the distribution of year-over-year annual earnings changes for workers in the ADP data between 2015 and 2016. We limit attention to workers who appear in 2015 and 2016 for the same number of months, plus or minus 3. The blue line plots the realized distribution in the ADP dataset, 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.

change of 0.01, in line with Guvenen et al. (2015). The standard deviation of annual earnings changes is 0.53 (compared with 0.51 in the SSA), while the skewness is -1.03 (vs -1.07) and kurtosis is 17.51 (vs 14.93).[4] We interpret these small differences to be the result of imperfectly capturing the annual earnings changes for job-changers and transitions to unemployment. Overall, however, we find the similarity of this figure to that in Guvenen et al. (2015) encouraging.

Appendix C Calculating Compensation Measures

This section details our construction of relevant compensation measures.

Appendix C.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

[4]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.

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. 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 between hourly statuses; however, it is worth bearing in mind when we present statistics by employee wage percentile.

Appendix C.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 are 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.[5]

To do so, we restrict our overtime imputation to those workers paid hourly. Given that the hours variables are essentially always set to 40 hours per week for salaried workers, we cannot separately distinguish overtime payments from bonuses and commissions for salaried workers. This implies that, by definition, we will have no overtime measures for salaried workers. For hourly workers, we infer overtime premiums implied using the hours and earnings subcategories. Specifically, 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

[5]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. However, the composition of earnings type across these five earnings categories is measured with error.

Figure A3: Distribution of Overtime Premiums

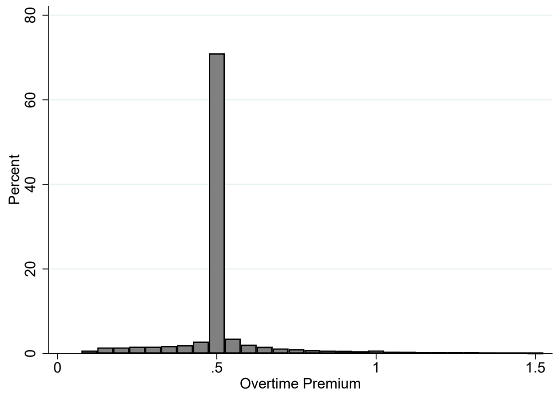

Notes: Figure shows the distribution of imputed overtime premiums for our employee sample for those with imputed overtime premiums greater than 1.1.

wage. Overtime wage rates, conversely, have large mass points at 1.5 times contract wages: about 80% of hourly workers with overtime premiums above 1.1 have implied wage rates which are 1.45-1.55 times their base wage with almost all of them exactly equal to 1.5. The distribution of our imputed overtime premium is shown in Appendix Figure A3.

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, we declare earnings in the overtime subcategory to be true overtime earnings for any worker whose imputed overtime wage is between 1.45 and 1.55 times their base earnings. There is no clear way to classify the remaining 20 percent of individuals with imputed overtime wages between 1.1 and 1.45 times their contract wage or those with overtime wages above 1.55 times their contract wage. As our base methodology, we include such earnings in our bonus measure. However, we also explored excluding those individuals from our sample all together when discussing the composition of compensation and cyclicality of compensation in Sections 3 and 4. Our key results were essentially identical under the two methods.

Appendix Figure A4 explores the heterogeneity in overtime hours worked more extensively. Focusing on those with positive overtime hours, Panel A shows that about one-quarter of those hourly workers with positive overtime hours work only between 1 and 5 hours of overtime during the month and just under half work less than 10 overtime hours. The median individual with positive overtime hours during the month is working about 11 extra hours during the month. However, there is a long right tail of overtime hours with about twenty percent of workers accruing over 40 additional overtime hours during the month. Panel B shows that 45 percent of all hourly workers who work overtime during a given year

Figure A4: Distribution of Overtime Hours for those With Positive Overtime Hours

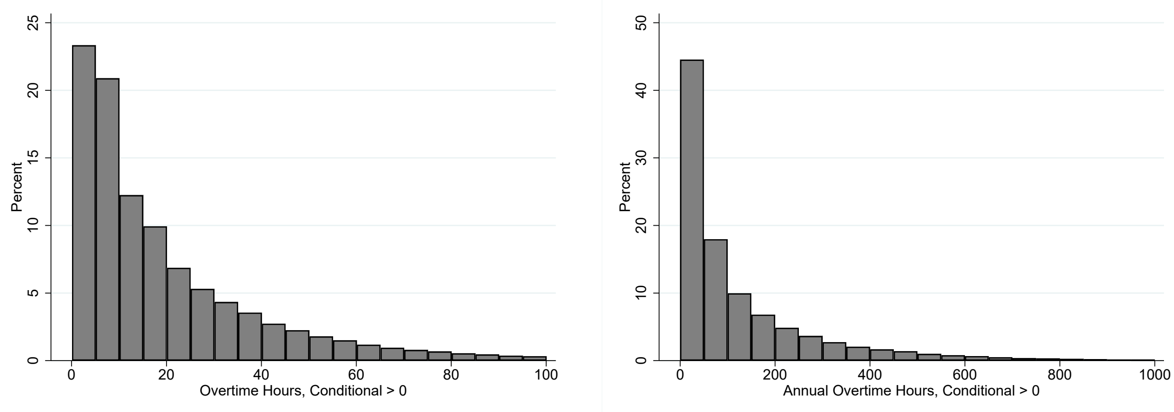

Panel A: Monthly Overtime Hours Distribution

Panel B: Annual Overtime Hours Distribution

Notes: Figure shows the distribution of monthly hours of overtime worked (Panel A) and annual hours of overtime worked (Panel B) conditional on overtime hours being positive during the respective time periods. Both panels restrict the sample to only those workers paid hourly whose overtime premium is approximately 1.5 times their contract wage. See text for additional details. Panels B further restricts the sample to those who remain continuously employed with the same employer for 12 consecutive calendar months.

work less than 50 annual overtime hours. Again, most annual overtime recipients work very little overtime during the year. This is consistent with the low overtime share of earnings for hourly workers highlighted in Table **??** and Figure **??** of the main text. But, as with the monthly distribution, there is a long tail of overtime hours with one-quarter of annual overtime recipients working over 200 annual overtime hours per year. Additionally, overtime earnings are a small component of annual earnings. Half of all overtime recipients receive less than 2 percent of their annual earnings from overtime. Only about 5 percent of overtime recipients receive more than 10 percent of their annual earnings from overtime compensation.

Appendix D Robustness of Nominal Wage Adjustments for Job-Stayers

In this section, we present robustness exercises for our key results - the distribution of nominal wage adjustments for job-stayers.

Appendix D.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 A5 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 **??**.

Figure A6 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 A6 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 **??** of the main text and Figure A5. 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 A5: 12-month Changes in December Base Wages, 24-month Job-Stayers

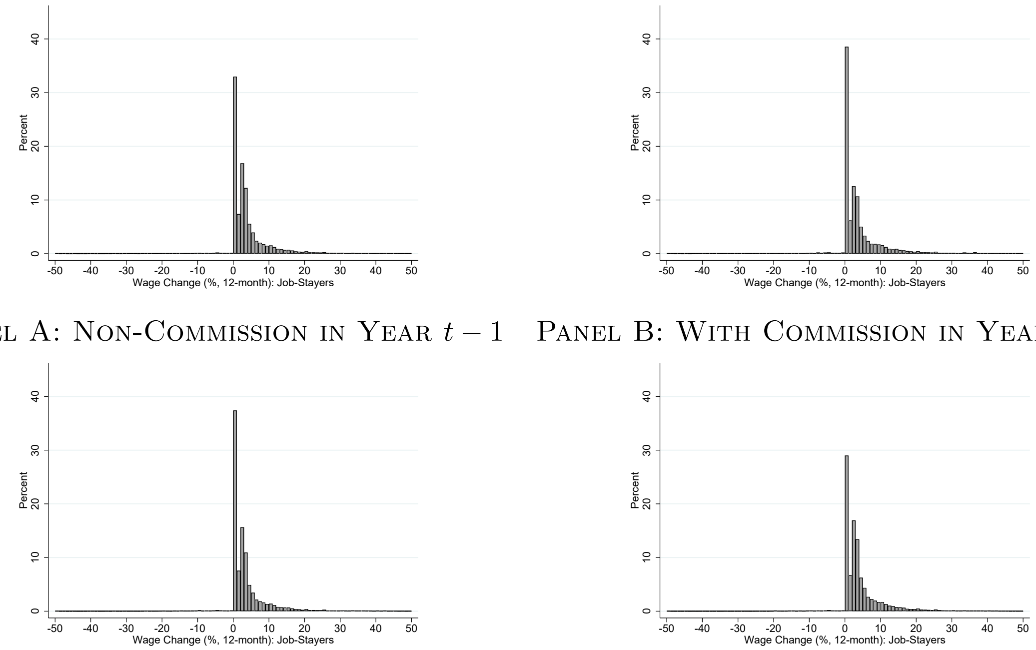

Panel A: Non-Commission in Year $t-1$ Panel B: With Commission in Year $t-1$

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 for 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.

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

Table A3 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 D.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 × 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: 12-month Base Wage Changes, Job-Stayers, Hourly Workers w/ Variable Hours

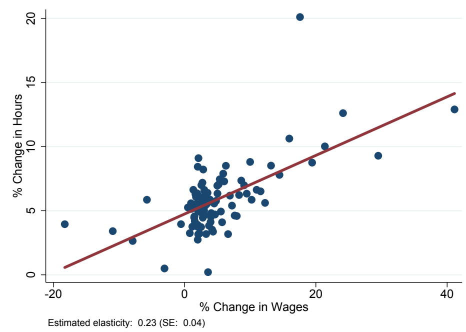

PANEL A: CHANGES IN HOURS VS CHANGES IN WAGES

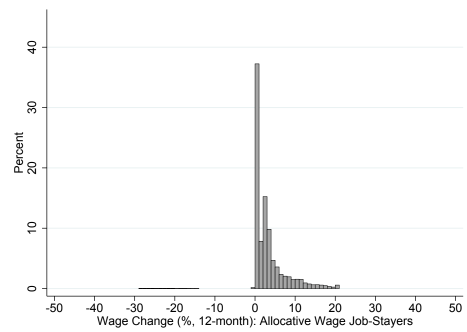

PANEL B: BASE WAGE CHANGE DISTRIBUTION

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 A3: Base Wage Change Statistics, Pooled 2008-2016 Sample of Job-Stayers

	Monthly	Quarterly	Annual
Unconditional			
Skewness of Wage Changes (%)	9.4	5.0	2.7
Kurtosis of Wage Changes (%)	155.8	42.8	13.3
Conditional on Any Wage Change			
Skewness of Wage Changes (%)	2.0	1.9	2.2
Kurtosis of Wage Changes (%)	13.1	11.1	10.5

Notes: 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.

Figure A7: 12-month base wage change distributions for job-stayers: unweighted

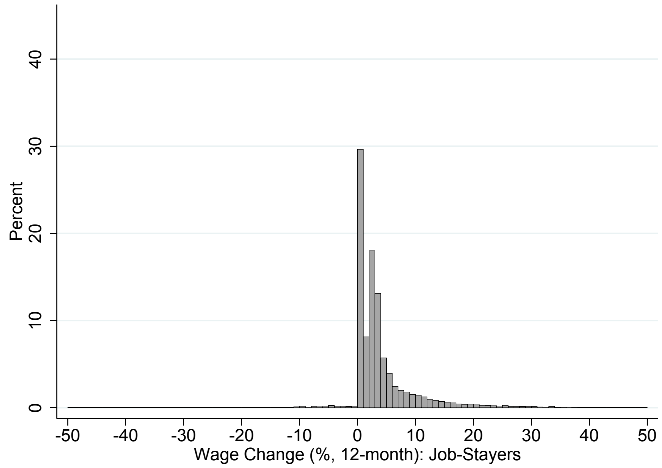

Notes: Figure plots the unweighted distribution of 12-month base wage changes for our employee sample of job-stayers. Figure is analogous to Figure **??** of the main text except pooling over hourly and salaried workers and not weighting to match the BLS's firm-size by industry distribution.

Figure A7 plots the unweighted distribution of 12-month base wage changes for job-stayers. The only difference between this figure and Figure **??** of the main text is that here we do not weight data in order to match the firm size × 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 × industry mix is fairly representative of the US economy, and 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 D.4 Mean Base Wage Change Size by Time Since Last Base Wage Change

Figure **??** from the main text provides 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 often have predictions regarding the average size of wage changes. These models predict that 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

Figure A8: Number of Nominal Base Wage Changes over a Calendar Year, Job-Stayer Sample

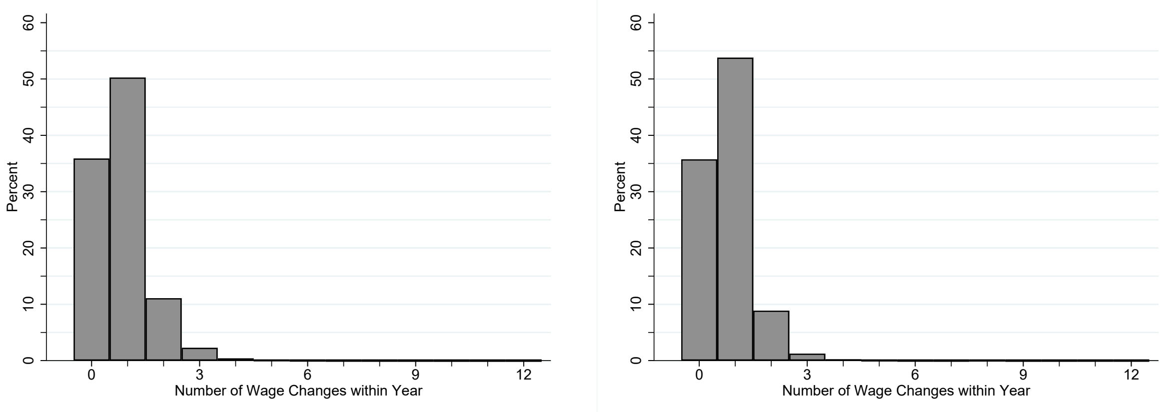

Panel A: Hourly Workers   Panel B: Salaried Workers

Note: Figure shows the distribution of the number of nominal base wage changes for hourly workers (left panel) and salaried workers (right panel) during a calendar year. 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 2016.

who wait longer between wage changes should observe larger average changes in absolute value. We explore this prediction next.

Appendix Figure A8 plots the histogram of number of base wage changes during a given calendar year for workers in our full-year employee sample. 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. Therefore, roughly 85 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.

Figure A9 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 A9 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, if a worker is promoted internally and the promotions are distributed throughout the year,

Figure A9: Mean Size of Base Wage Changes by Time Since Last Change, Job-Stayers

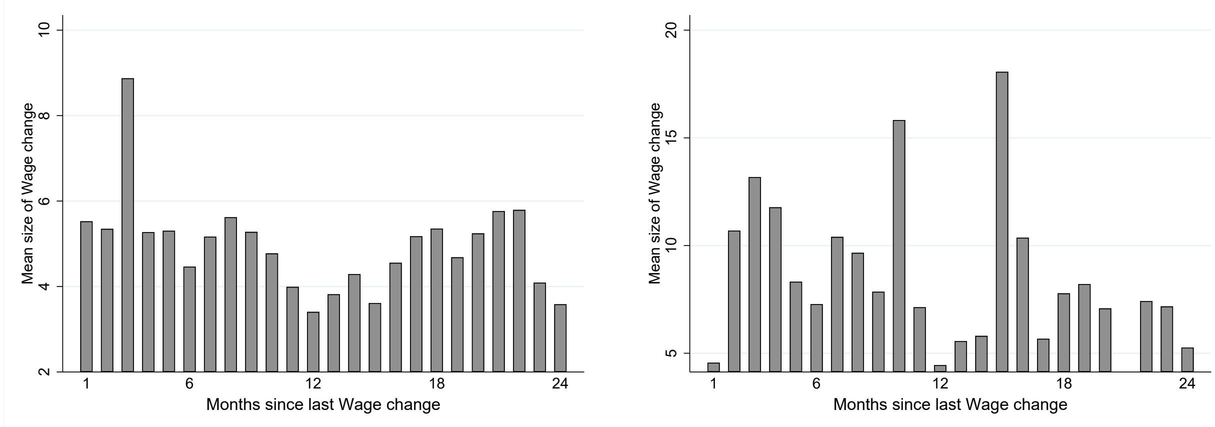

Panel A: Hourly Workers　　　Panel B: Salaried Workers

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.

workers who receive a base wage change off cycle would also get larger base wage changes.

Appendix E 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.

Appendix Figure A10 shows the probability 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, 64.4% of hourly workers and 66.9% of salaried workers in firms with under 500 employees receive a base wage change. The comparable numbers for firms with 5000+ employees are 80.2% and 77.1%, 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 explain why our weighted results and unweighted results are so similar to each other.

Appendix Figure A10 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 **??**. 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. Appendix Figure A11 plots the distribution of 12-month base wage changes for job-stayers in this small firm sample for the period 2014-2016.[6] 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 2.5%. 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 A11 to compute a new measure of the probability of nominal base wage adjustments inclusive of workers in small firms. Accounting for small firms leads to a corrected annual probability of base wage change of 60.7% ($0.271 \times 51.3\% + (1 - 0.271) \times 64.3\%$), while the probability of a year-over-year base wage cut is 2.4% for all

[6]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 measures 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 A10: Share with Base Wage Change by Firm Size and Industry, All Years

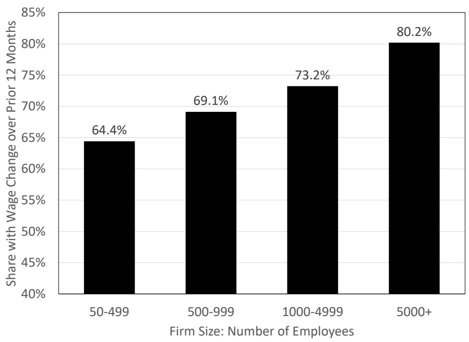

Panel A: Hourly Workers by Size

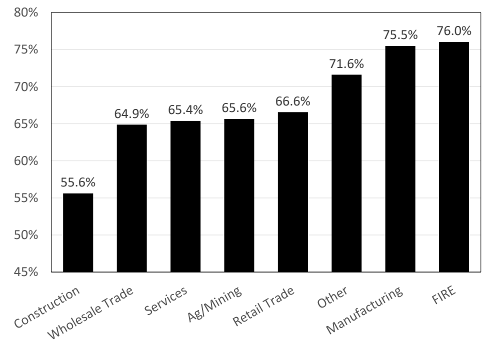

Panel B: Hourly Workers by Industry

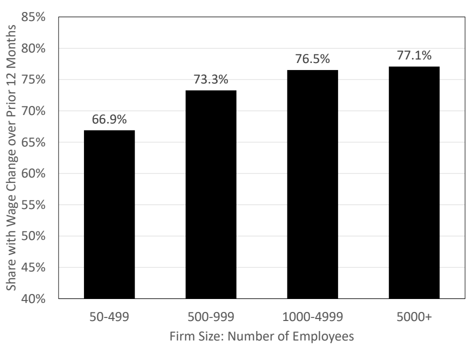

Panel C: Salaried Workers by Size

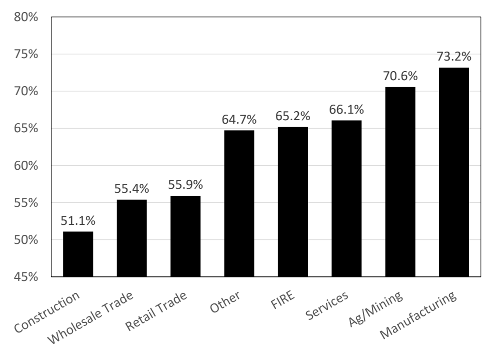

Panel D: Salaried Workers by Industry

Note: Figure shows the probability of receiving a base wage change over a 12-month period 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.

workers inclusive of those at small firms. Note, that these adjusted probabilities are very close to those reported in the main text including only workers in firms with more than 50 employees (64.3% vs. 60.7% and 2.5% vs. 2.4%). We conclude that the omission of workers in firms with less than 50 employees does not bias our results substantively.

Figure A11: 12-month base wage change distributions for job-stayers: firms with less than 50 employees

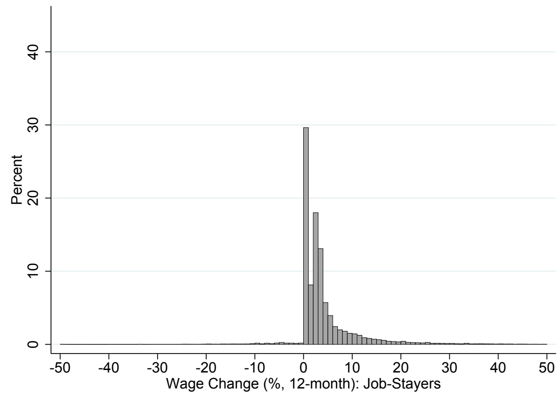

Panel A: $\geq$ 50 Employee Firms - Unweighted

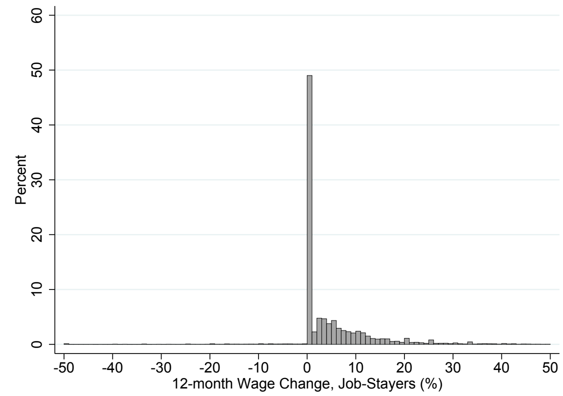

Panel B: $<$ 50 Employee Firms - Unweighted

Notes: Left panel of the figure plots the 12-month base wage change for job-stayers for a sample of workers employed in firms with more than 50 employees. For this figure, we show the raw unweighted data. Right panel of the 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 also otherwise unweighted.

Appendix F Robustness of Cyclicality Regressions

In the top panel of Appendix Table A4 we ask whether the intensive margin of overtime and bonuses varies with changes in state unemployment. The regression in Panel A is the same as the cyclicality regression in the main text except that y_{ijst} is the log of overtime hours (columns 1 and 2 for those with positive overtime hours) and the log of bonus payments (columns 3 and 4for those with positive bonuses). Columns 1 and 3 exclude individual fixed effects while columns 2 and 4 include individual fixed effects. Regardless of individual fixed effects, the amount of overtime hours (conditional on overtime hours existing) and the size of bonuses (conditional on bonuses being positive) are relatively acyclical.

The bottom panel of Appendix Table A4 explores the cyclicality of overtime receipt across industries. The first two columns restrict the sample to those workers in manufacturing industries, while the second two columns consider those in non-manufacturing industries. For each industry subset, we highlight results without and with individual fixed effects. The dependent variable for these regressions is the propensity to receive overtime hours. Otherwise, the regressions are the same as in the top panel. We split the sample by industry since the BLS tracks annual manufacturing hours. The BLS data highlight that overtime hours within the manufacturing sector are pro-cyclical. The BLS data do not control for

Table A4: Robustness of Cyclicality of Various Forms of Compensation

Panel A: Cyclicality of Compensation Components

	Log Overtime (1)	Log Overtime (2)	Log Bonus (3)	Log Bonus (4)
Δ Unemployment Rate (%)	-0.04	-0.02	-0.02	-0.01
	(0.01)	(0.00)	(0.01)	(0.00)
Workers Included	Hourly	Hourly	All	All
State and Industry FE	Y	Y	Y	Y
Individual FE	N	Y	N	Y
Observations (000s)	273	273	348	348
Mean of Dep. Var.	3.95	3.95	7.95	7.95

Panel B: Cyclicality of Overtime for Manufacturing and Non-Manufacturing Sectors

	Manufacturing		Non-Manufacturing	
	% With Overtime (1)	% With Overtime (2)	% With Overtime (3)	% With Overtime (4)
Δ Unemployment Rate (%)	-0.82	-0.54	-0.54	-0.19
	(0.23)	(0.13)	(0.29)	(0.13)
Workers Included	Hourly	Hourly	Hourly	Hourly
State and Industry FE	Y	Y	Y	Y
Individual FE	N	Y	N	Y
Observations (000s)	94	94	320	320
Mean of Dep. Var.	74.5	74.5	59.9	59.9

Notes: Table reports robustness specifications for Table **??** of the main text. The top panel is the same as the top panel of Table **??** of the main text except we exclude individual fixed effects from all regression columns. The bottom panel shows overtime receipt cyclicality separately for workers in the manufacturing and non-manufacturing sectors (with and without individual fixed effects). Otherwise, the regressions are the same as show in Table **??** of the main text. See main text for additional details.

individual fixed effects. As seen from columns 1 and 3 of Panel B, the propensity to receive overtime is slight more pro-cyclical in manufacturing than in non-manufacturing industries without controlling for individual fixed effects. However, as seen in column 2, controlling for individual fixed effects explains about a third of the observed pro-cyclicality of overtime hours within the manufacturing sector. Even with individual fixed effects, overtime remains statistically significantly pro-cyclical in the manufacturing sector with a coefficient of -0.54 and standard error of 0.13. However, within the non-manufacturing sectors – which make

Table A5: Annual Persistence of Base Wage, Bonuses, and Overtime, sample of full-year job-stayers, 2009-2016

	Log December Base Wage (1)	% With Bonus (2)	Log Annual Bonus (3)	% With Overtime (4)	Log Annual Overtime Hours (5)
y_{it-1}	0.82	0.57	-0.02	0.72	0.05
	(0.01)	(0.00)	(0.02)	(0.01)	(0.02)
Observations (000s)	463	463	115	307	235
Individual FE	Y	N	Y	N	Y
Workers Included	All	All	All	Hourly	Hourly

Notes: Table reports OLS-estimated AR(1) coefficients from appendix equation (A1). White heteroskedasticity robust standard errors clustered at the employee-level reported in parentheses. We use the employee sample restricting our analysis to workers who remain continuously employed with the same firm for two consecutive calendar years.

up the bulk of employment – there is no significant relationship between the propensity to work overtime and unemployment rates once controlling for individual fixed effects. These results highlight that even if overtime is procyclical in the manufacturing sector, the same patterns do not exist in other sectors after controlling for individual fixed effects.

Appendix G Persistence Regressions

In this section of the Online Appendix, we estimate autocorrelation coefficients for our various components of compensation. To do so, we take advantage of within worker wage dynamics within the same job. Specifically, we estimate OLS regressions of the form:

$$y_{it} = \rho y_{it-1} + \alpha_i + \epsilon_{it} \tag{A1}$$

where i indexes a worker-firm pair, t represents a year and α_i is a worker-job fixed effect. y_{it} represents the value of a particular compensation measure for job i in year t. For this exercise we use our sample of workers who remain continuously employed on the same job (job-stayers) for two consecutive calendar years, in order to measure annual adjustments in bonus and overtime compensation.

Appendix Table A5 reports the estimated autocorrelation coefficient ρ of various compensation measures at the individual level. We explore the persistence of log base wages (first column) and various specifications of overtime and bonus receipt (columns 2 through 5). As seen from the table, base wages are highly persistent with an annual autocorrelation of 0.82. Columns 2 and 4 show that that propensity to accrue overtime hours and to receive bonuses, respectively, are also highly persistent over time. If a worker receives a bonus this year, they are very likely to receive a bonus next year. Likewise, people who accrue overtime this year are also likely to accrue overtime next year. However, the amount of overtime compensation and the amount of bonus receipt, conditional on receipt, are essentially i.i.d.. Indeed regressing the share of pay in bonuses and overtime (not shown) on its lagged value yields a negative coefficient, suggesting that high bonus years tend to be followed by low bonus years. These results suggest that base wage adjustments may be a far better measure of permanent wage adjustments than are bonus payments. Given this, the ability to adjust base wages is likely more important for changes in the user cost of labor than the ability to adjust bonuses.

References

ADP, LLC, "ADP Confidential Payroll Files (2008-2016)." First Accessed: January, 2017.

Brown, Charles and James Medoff, "The Employer Size Wage Effect," *Journal of Political Economy*, 1989, *95* (5), 1027–1059.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren, "Integrated Public Use Microdata Series, Current Population Survey: Version 8.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D030.V8.0." Accessed: June 7, 2018.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," 2016.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?," 2015.

U.S. Census Bureau, "Business Dynamics Statistics: Legacy BDS Firm Characteristics Data Tables 1977-2014 (September, 2016). https://www.census.gov/programs-surveys/bds/data/data-tables/legacy-firm-characteristics-tables-1977-2014.html." Accessed: June 9, 2017.